High-Throughput Science

Hanspeter Pfister, Harvard University
pfister@seas.harvard.edu

Scientific Data Explosion

Astronomy

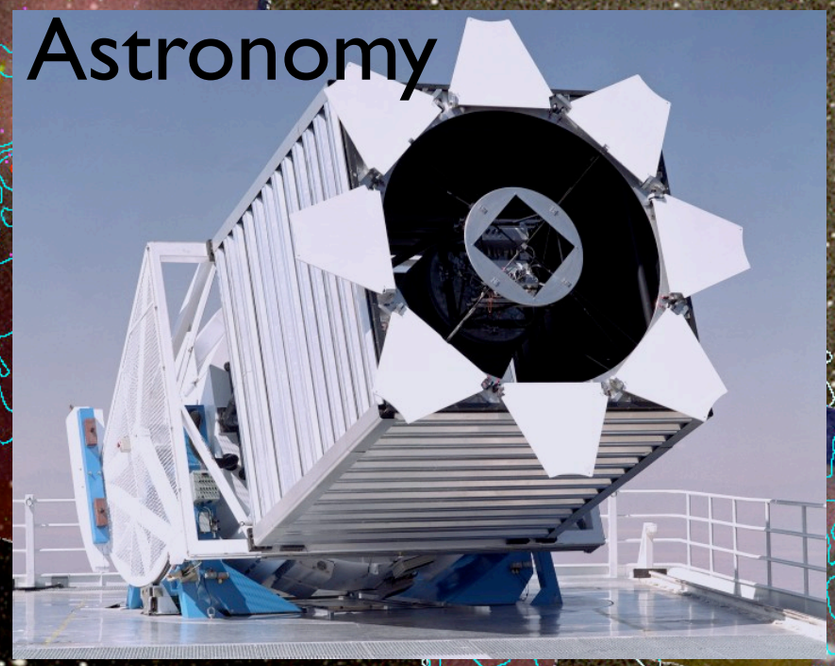

Microscopy

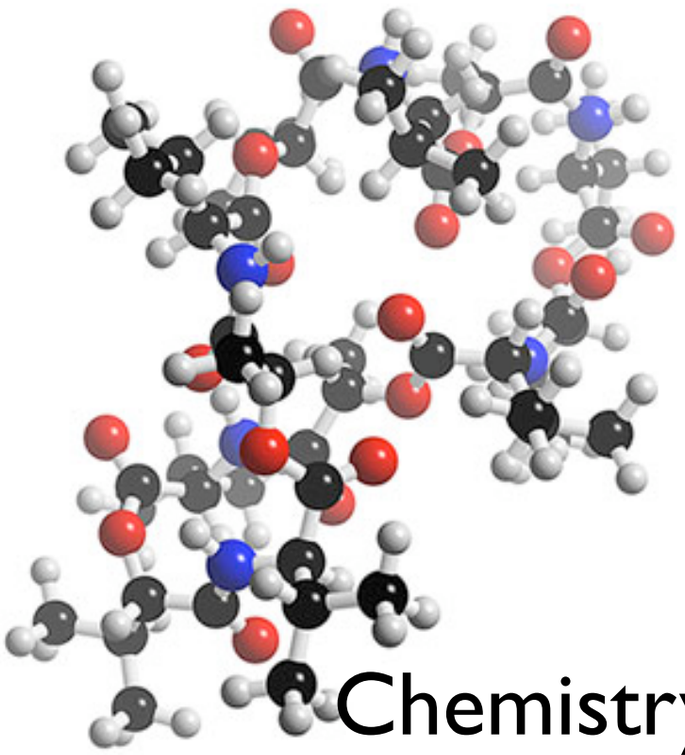

Chemistry

Social Sciences

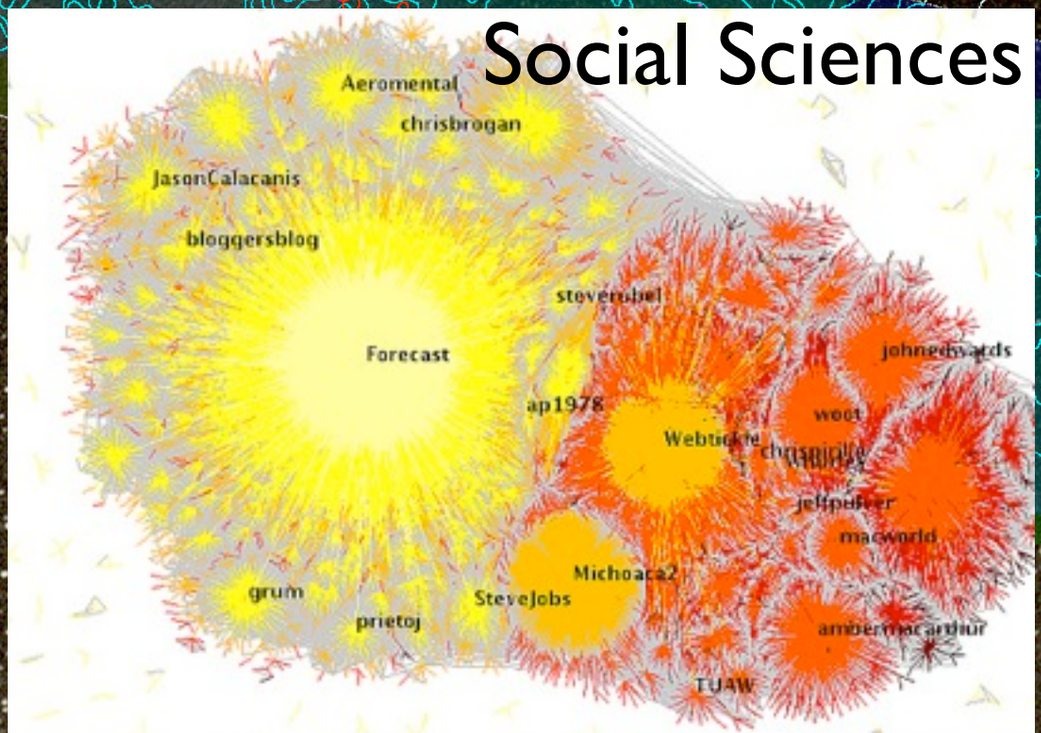

Twitter Social Network © UMBC eBiquity Research Group

Computing Revolution

- Commodity energy-efficient supercomputers

4 Teraflops @ 1400 W

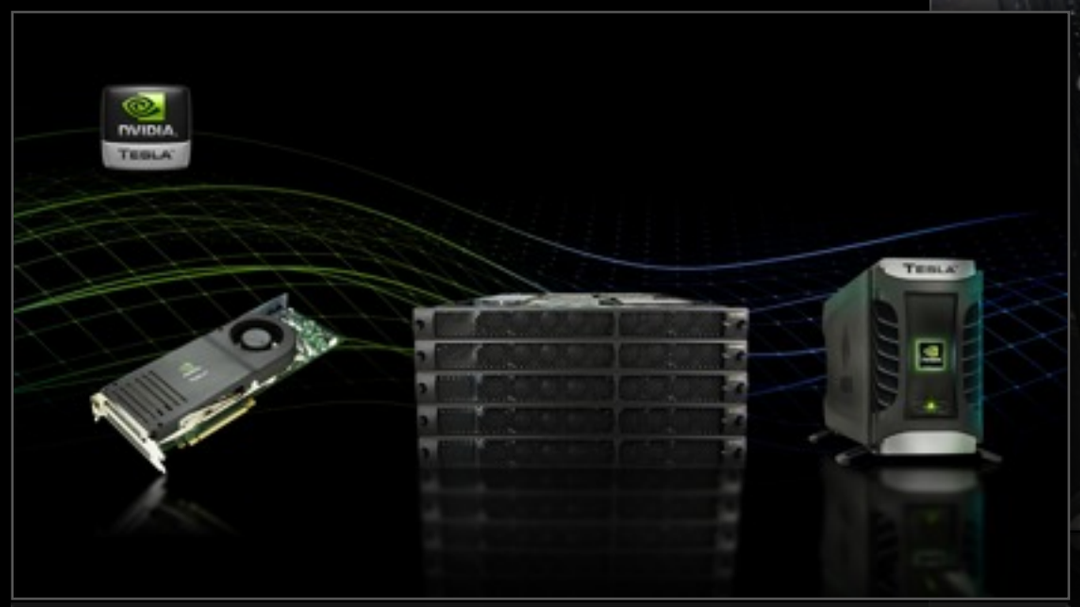

Traditional

Sensor

Modeling
Simulation
Analysis

Traditional

Sensor

Modeling
Simulation
Analysis

Challenge

- Limited by communication, not processor performance [Jim Grey, Distributed Computing Economics]

Challenge

- Limited by communication, not processor performance [Jim Grey, Distributed Computing Economics]

Challenge

Sensor

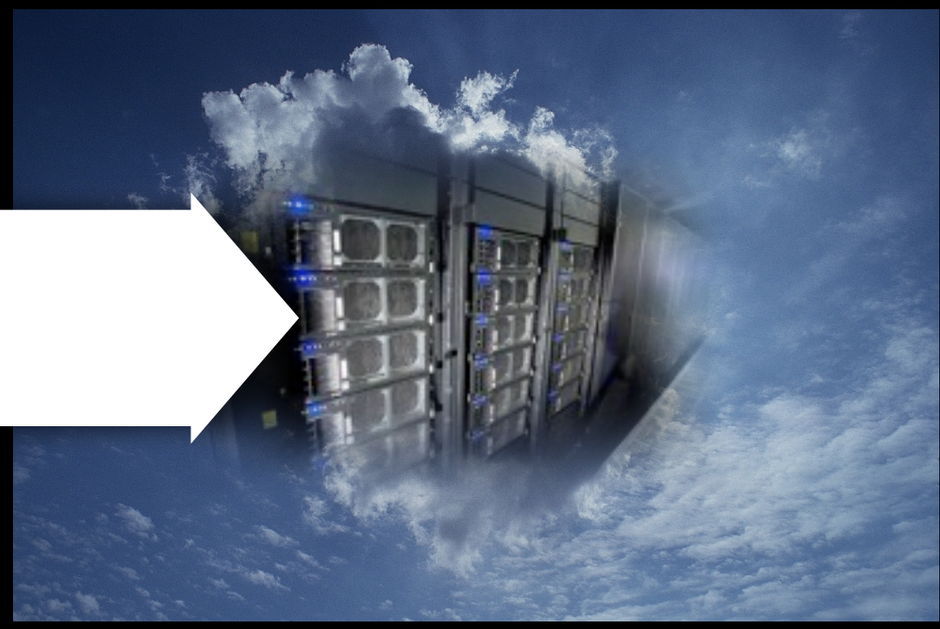

Modeling
Simulation
Analysis

Challenge

Sensor

Modeling
Simulation
Analysis

Challenge

Modeling
Simulation
Analysis

High-Throughput Computing

Sensor

Modeling
Simulation
Analysis

High-Throughput Computing

Sensor

HTC

Modeling
Simulation
Analysis

High-Throughput Computing

Sensor

HTC

Modeling
Simulation
Analysis

High-Throughput Computing

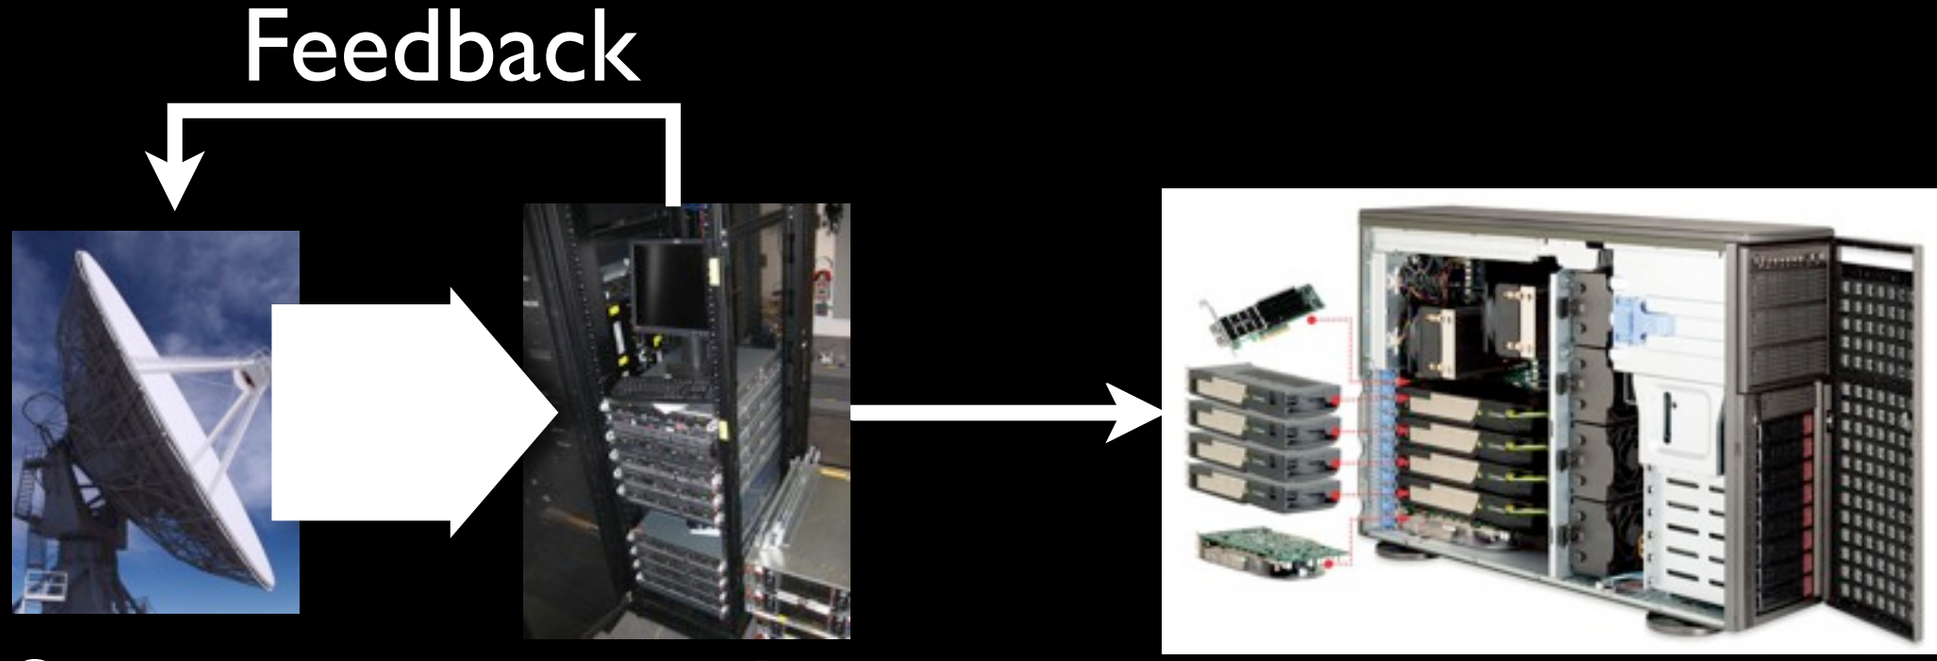

High-Throughput Computing

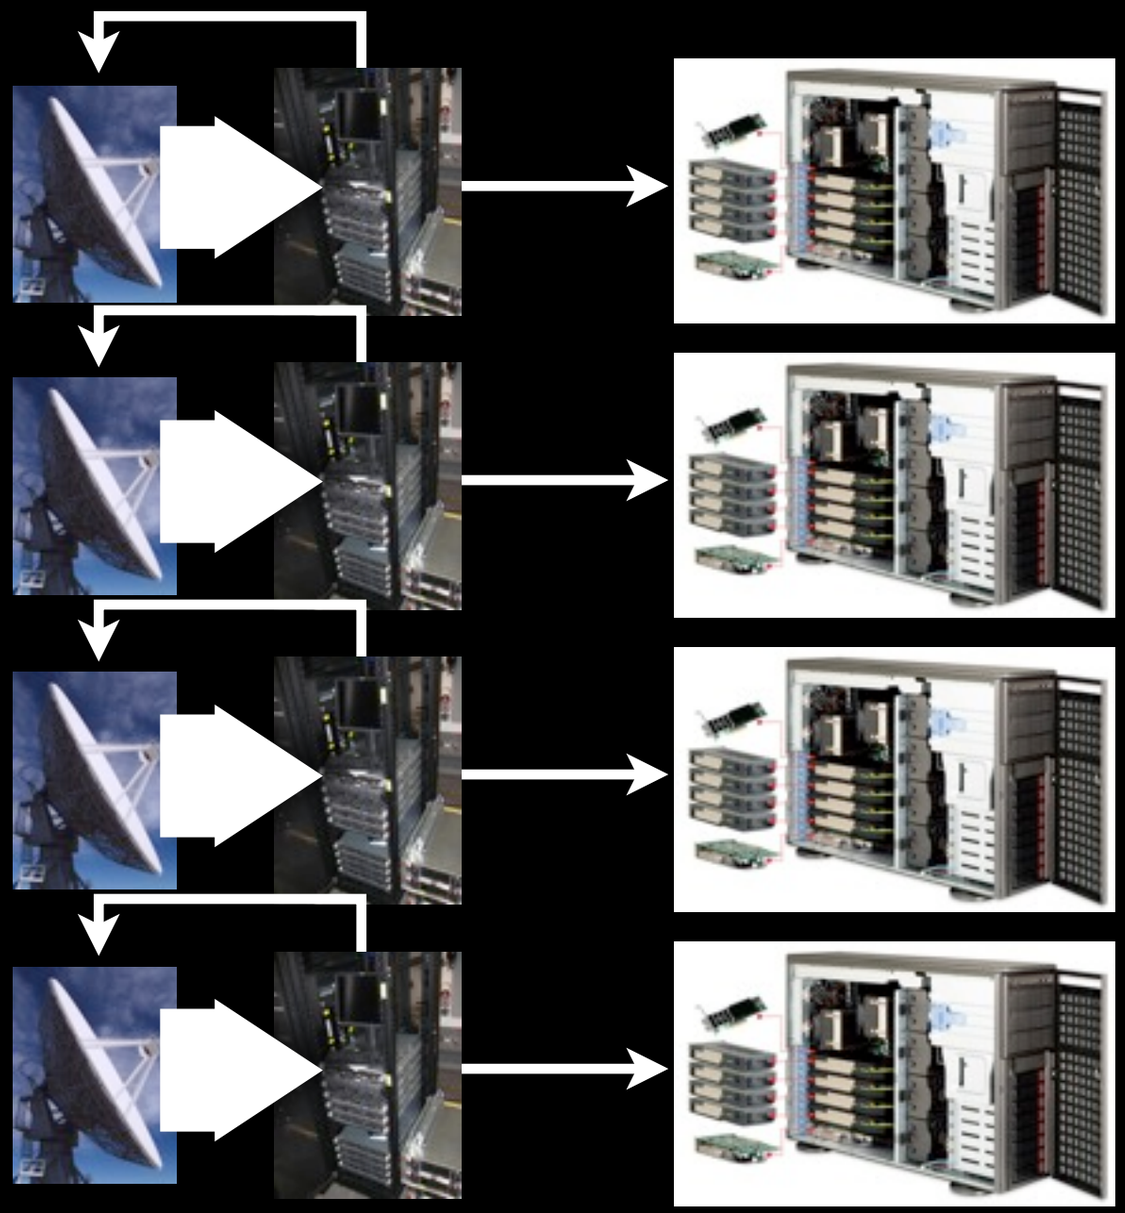

Outline

- How is the brain wired?
- How did the universe start?
- How does matter interact at the quantum level?
- How does the human visual system work?
- How can we prevent heart attacks?

The Connectome

Discovering the Wiring Diagram of the Brain

Collaborators

- Harvard Center for Brain Science
  - Jeff Lichtman & Clay Reid
- Kitware Inc.
  - Will Schroeder, Charles Law, Rusty Blue
- VRVis Vienna
  - Markus Hadwiger, Johanna Beyer
- SEAS
  - Amelio Vazquez, Won-Ki Jeong
  - Hanspeter Pfister



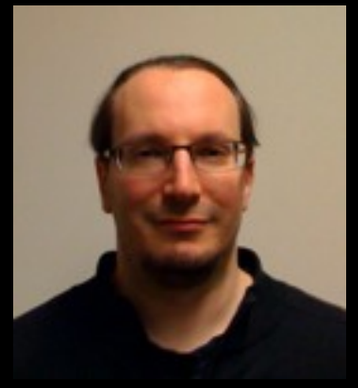


The Scientific Challenge

- What do large brain circuits do?

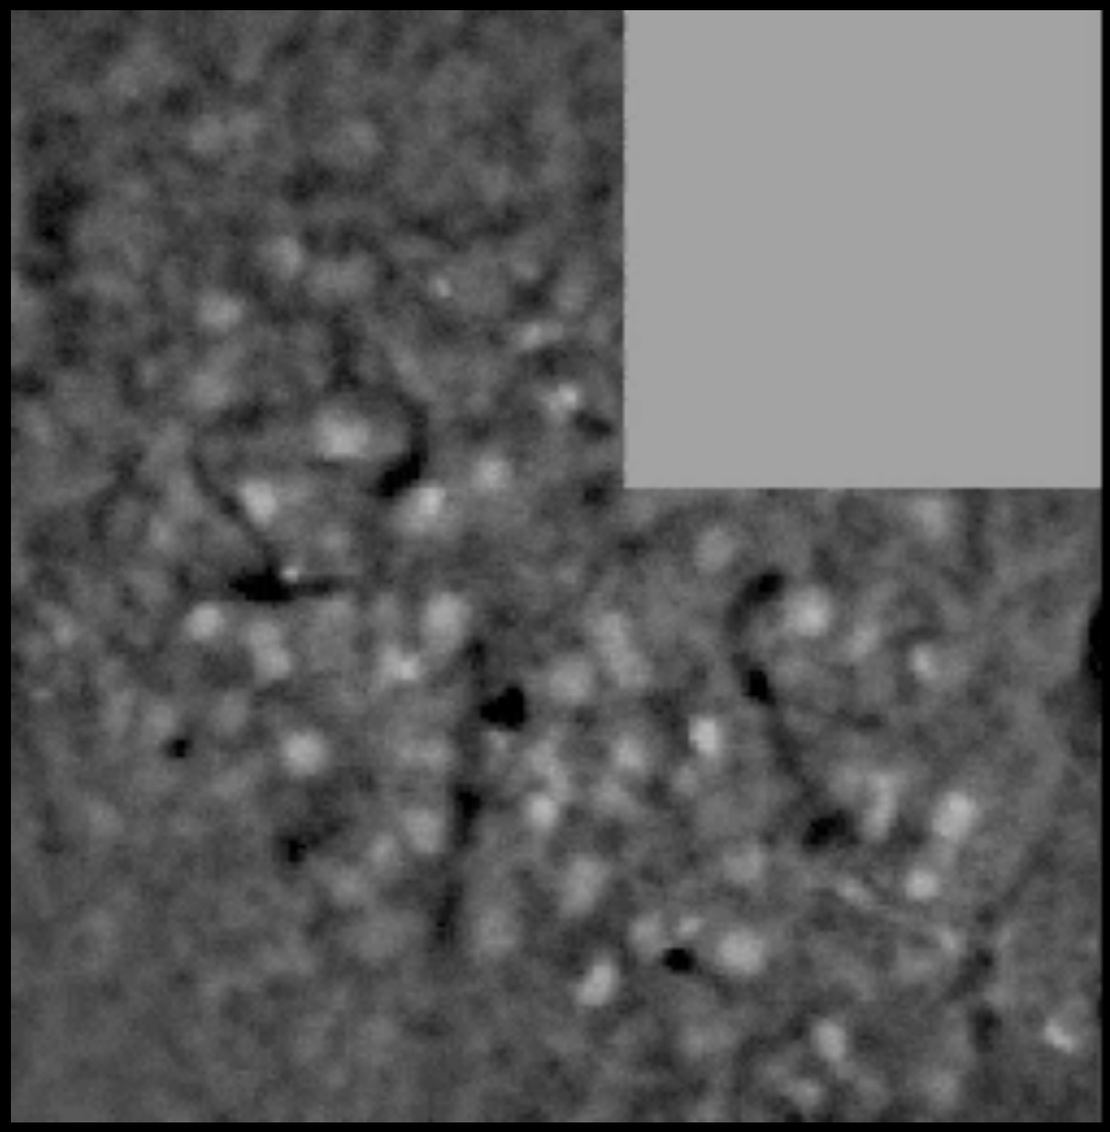

C. Reid

The Scientific Challenge

- What do large brain circuits do?

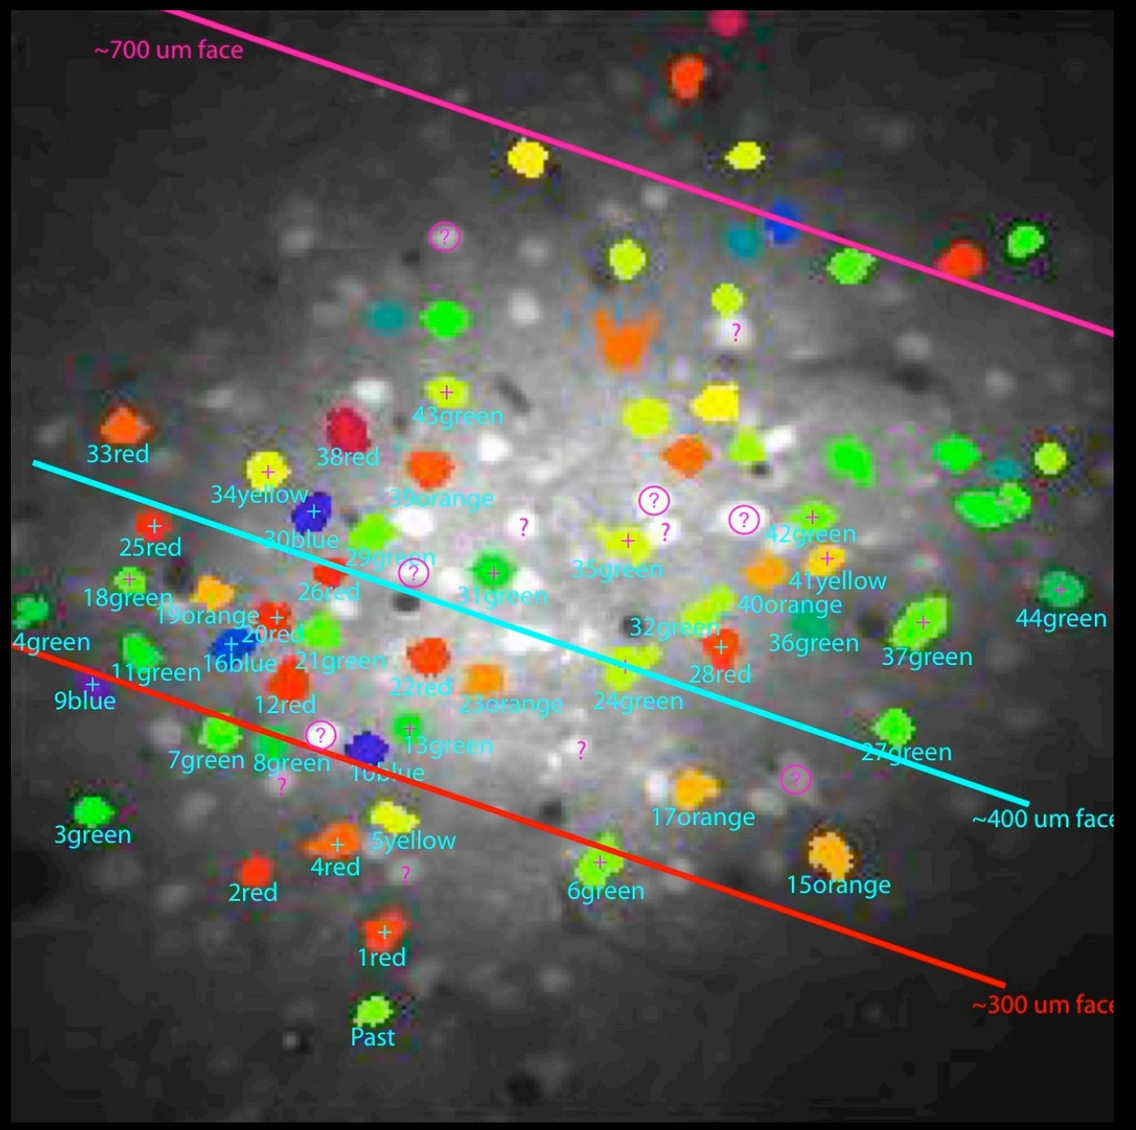

C. Reid

The Scientific Challenge

- What is their connectivity?

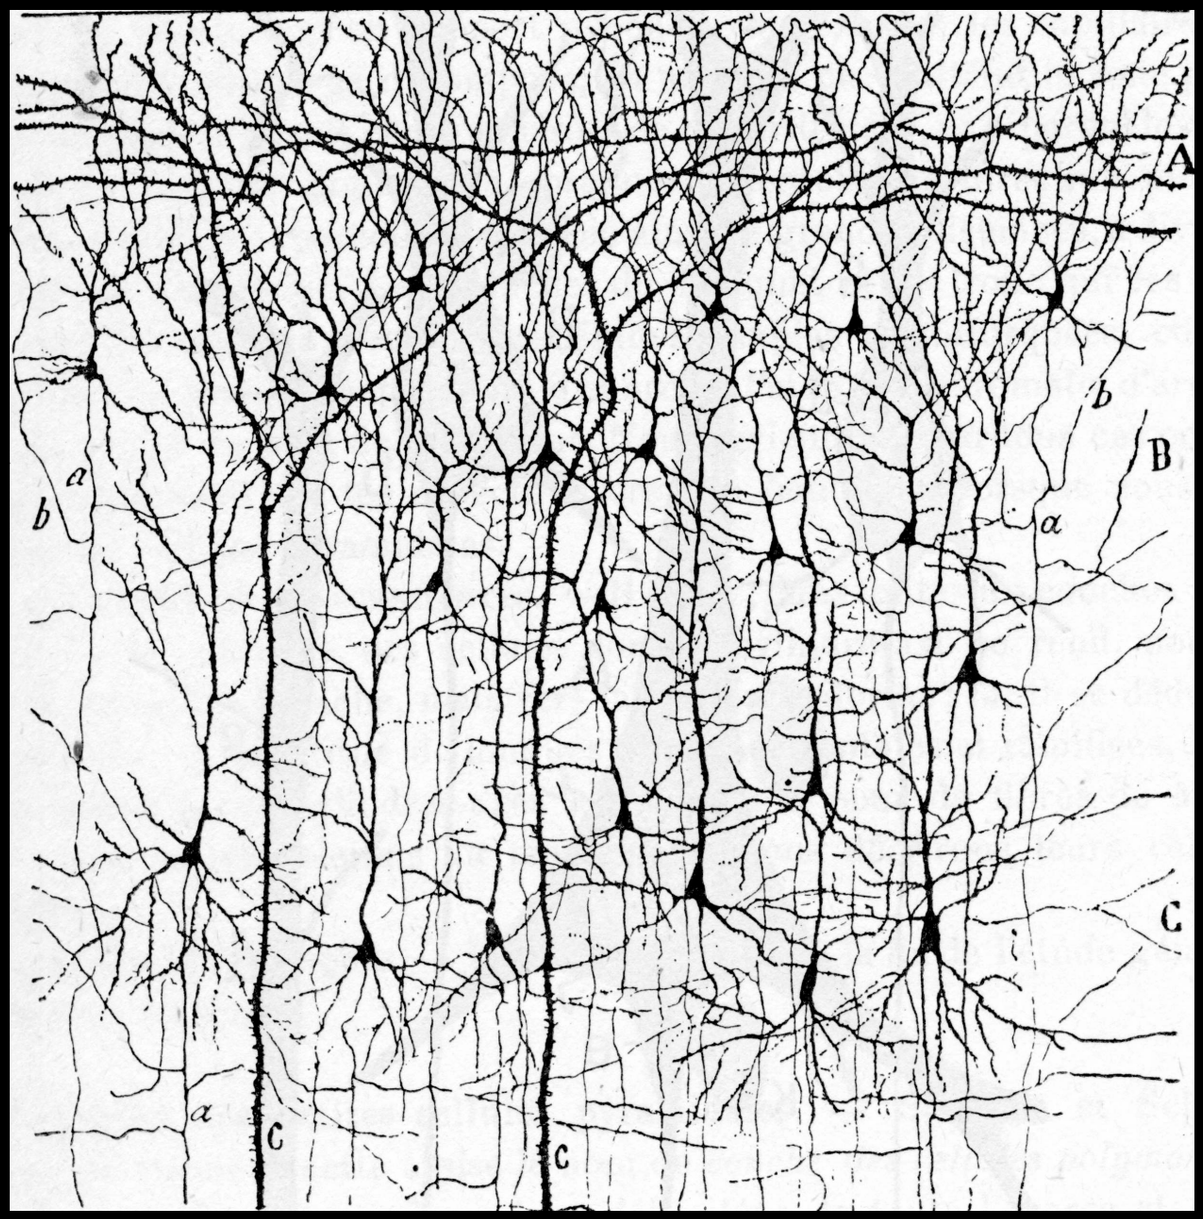

Ramón y Cajal, 1905

Connectome Pipeline

Connectome Pipeline

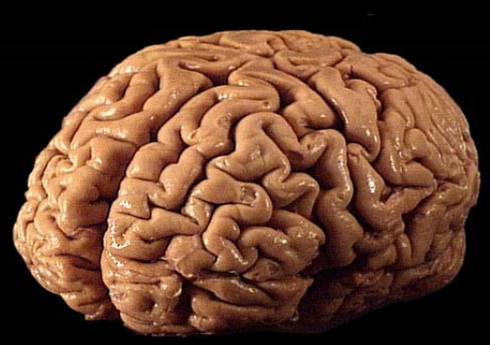

Get a brain

Connectome Pipeline

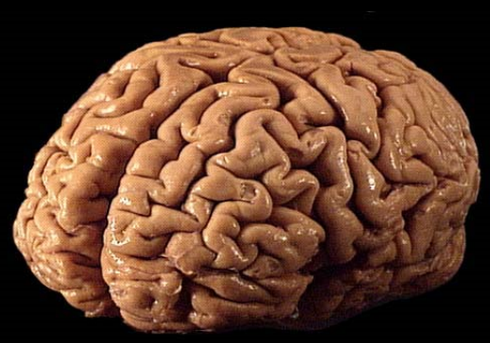

Get a brain

Get a piece

Connectome Pipeline

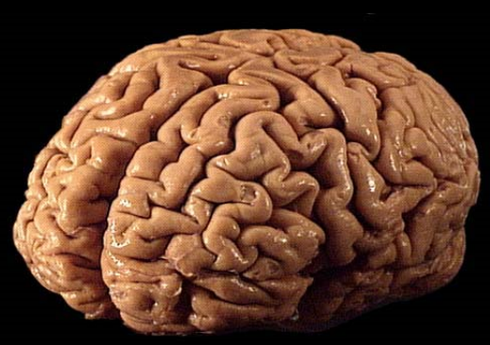

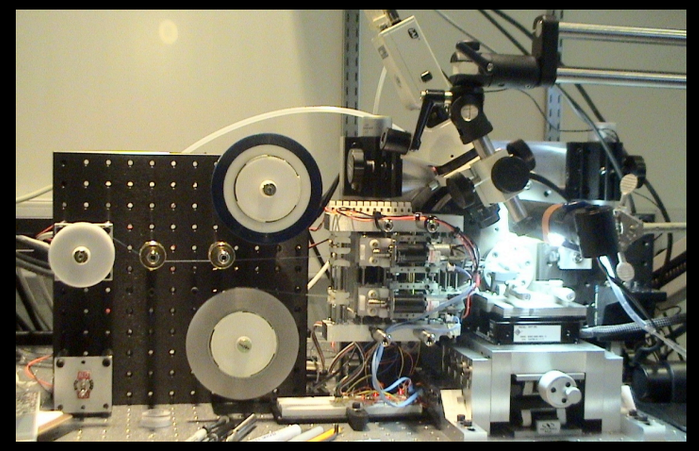

Get a brain → Get a piece → Cut it thin

Automatic Tape-collecting Lathe Ultra Microtome

K. Hayworth & B. Kasthuri

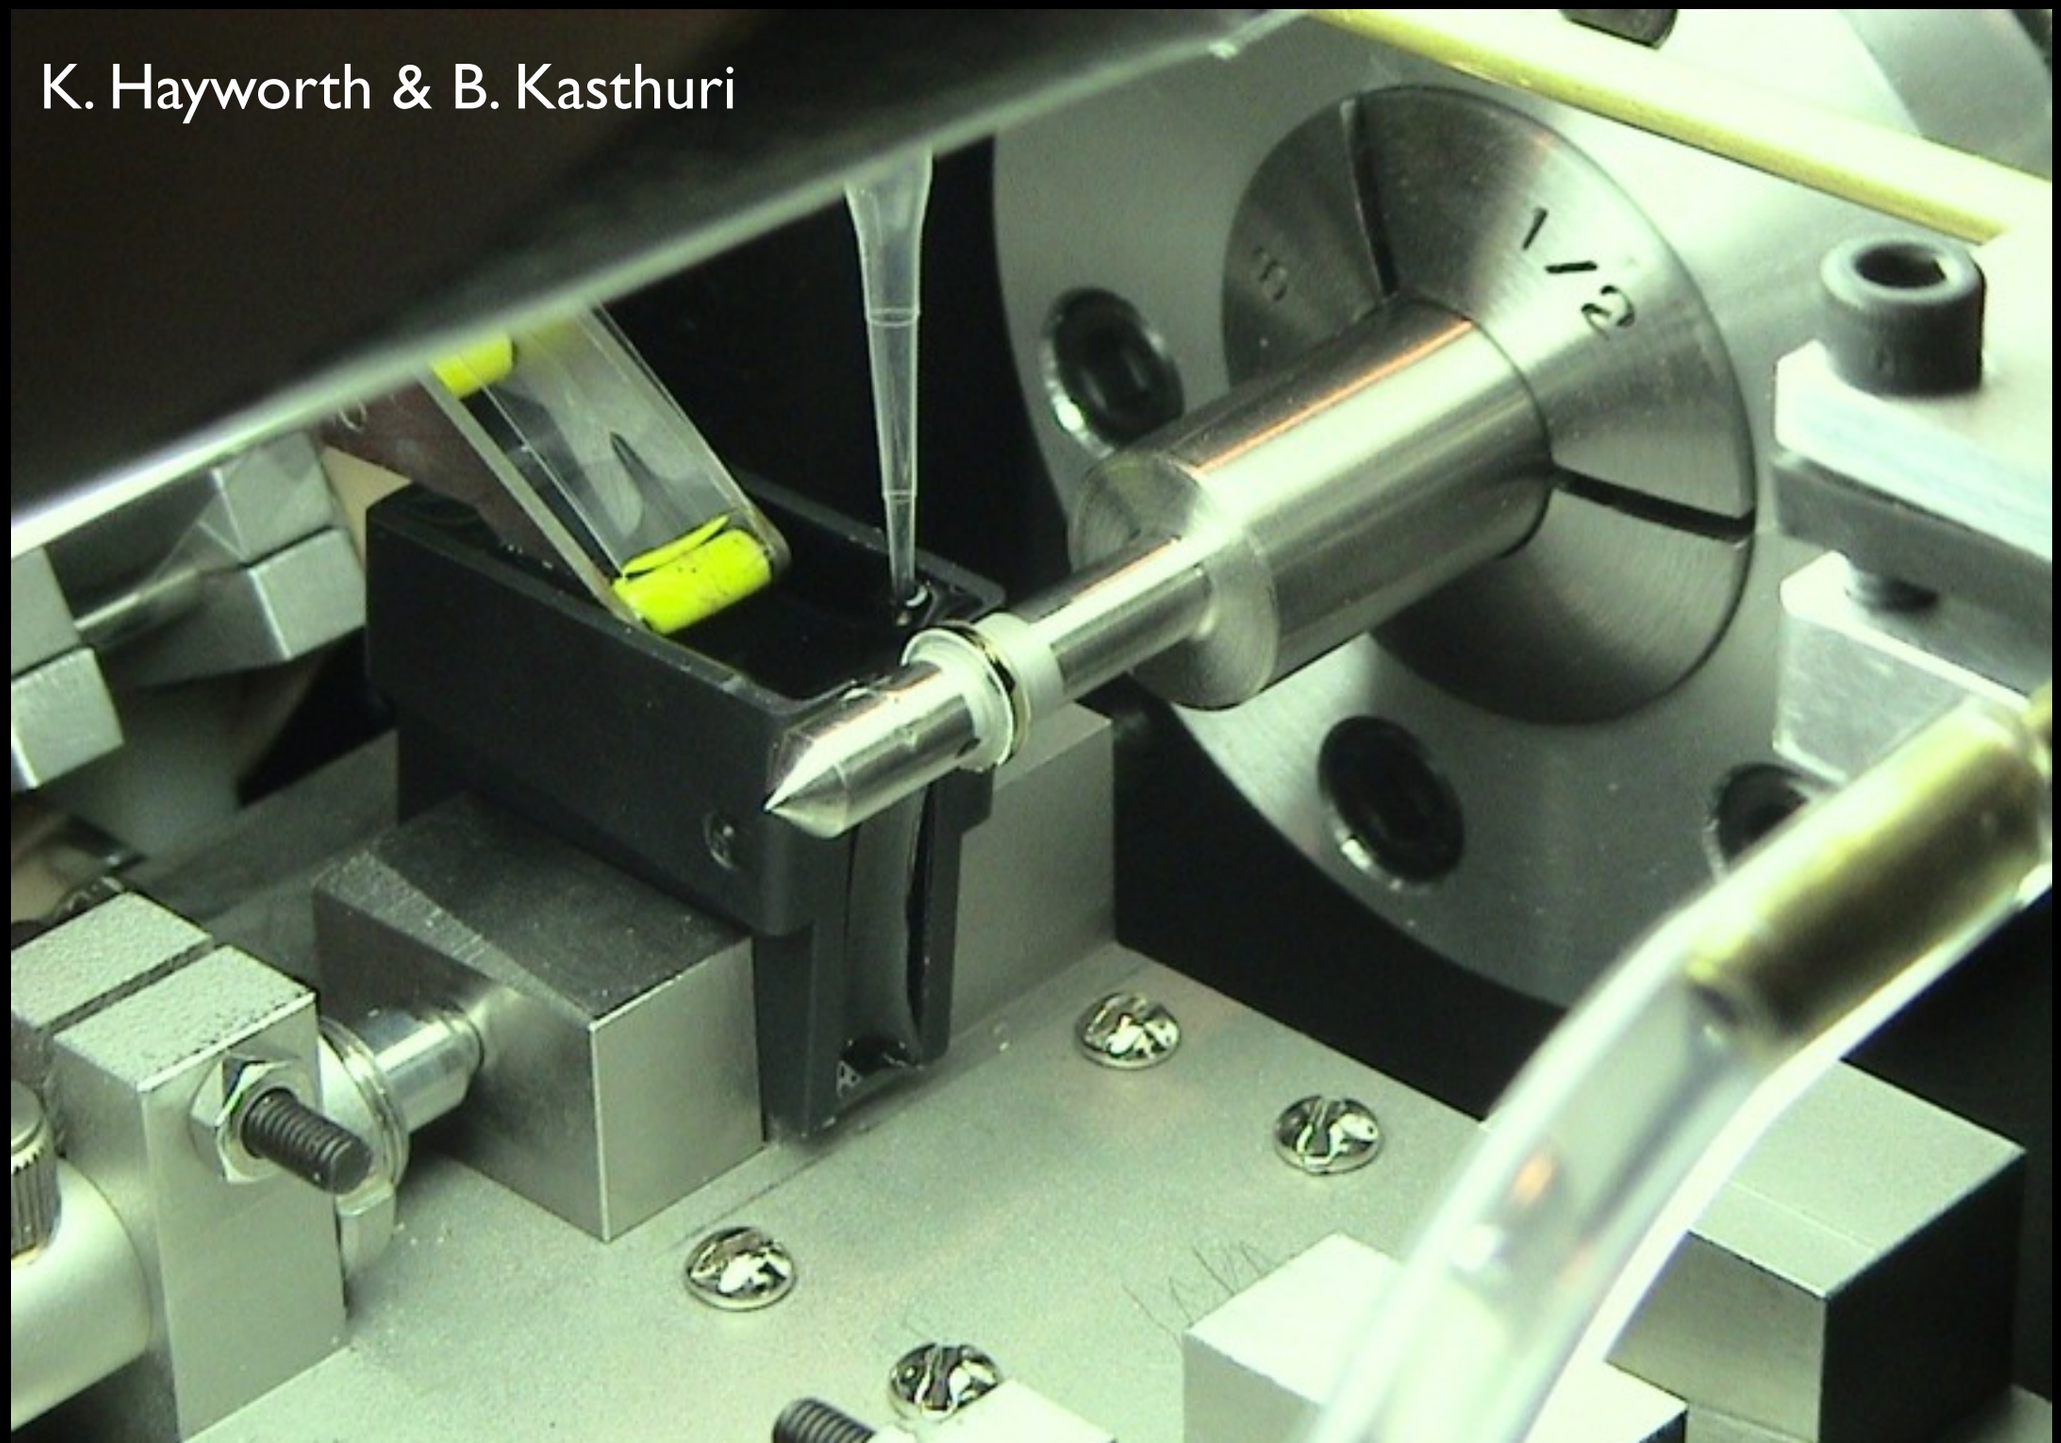

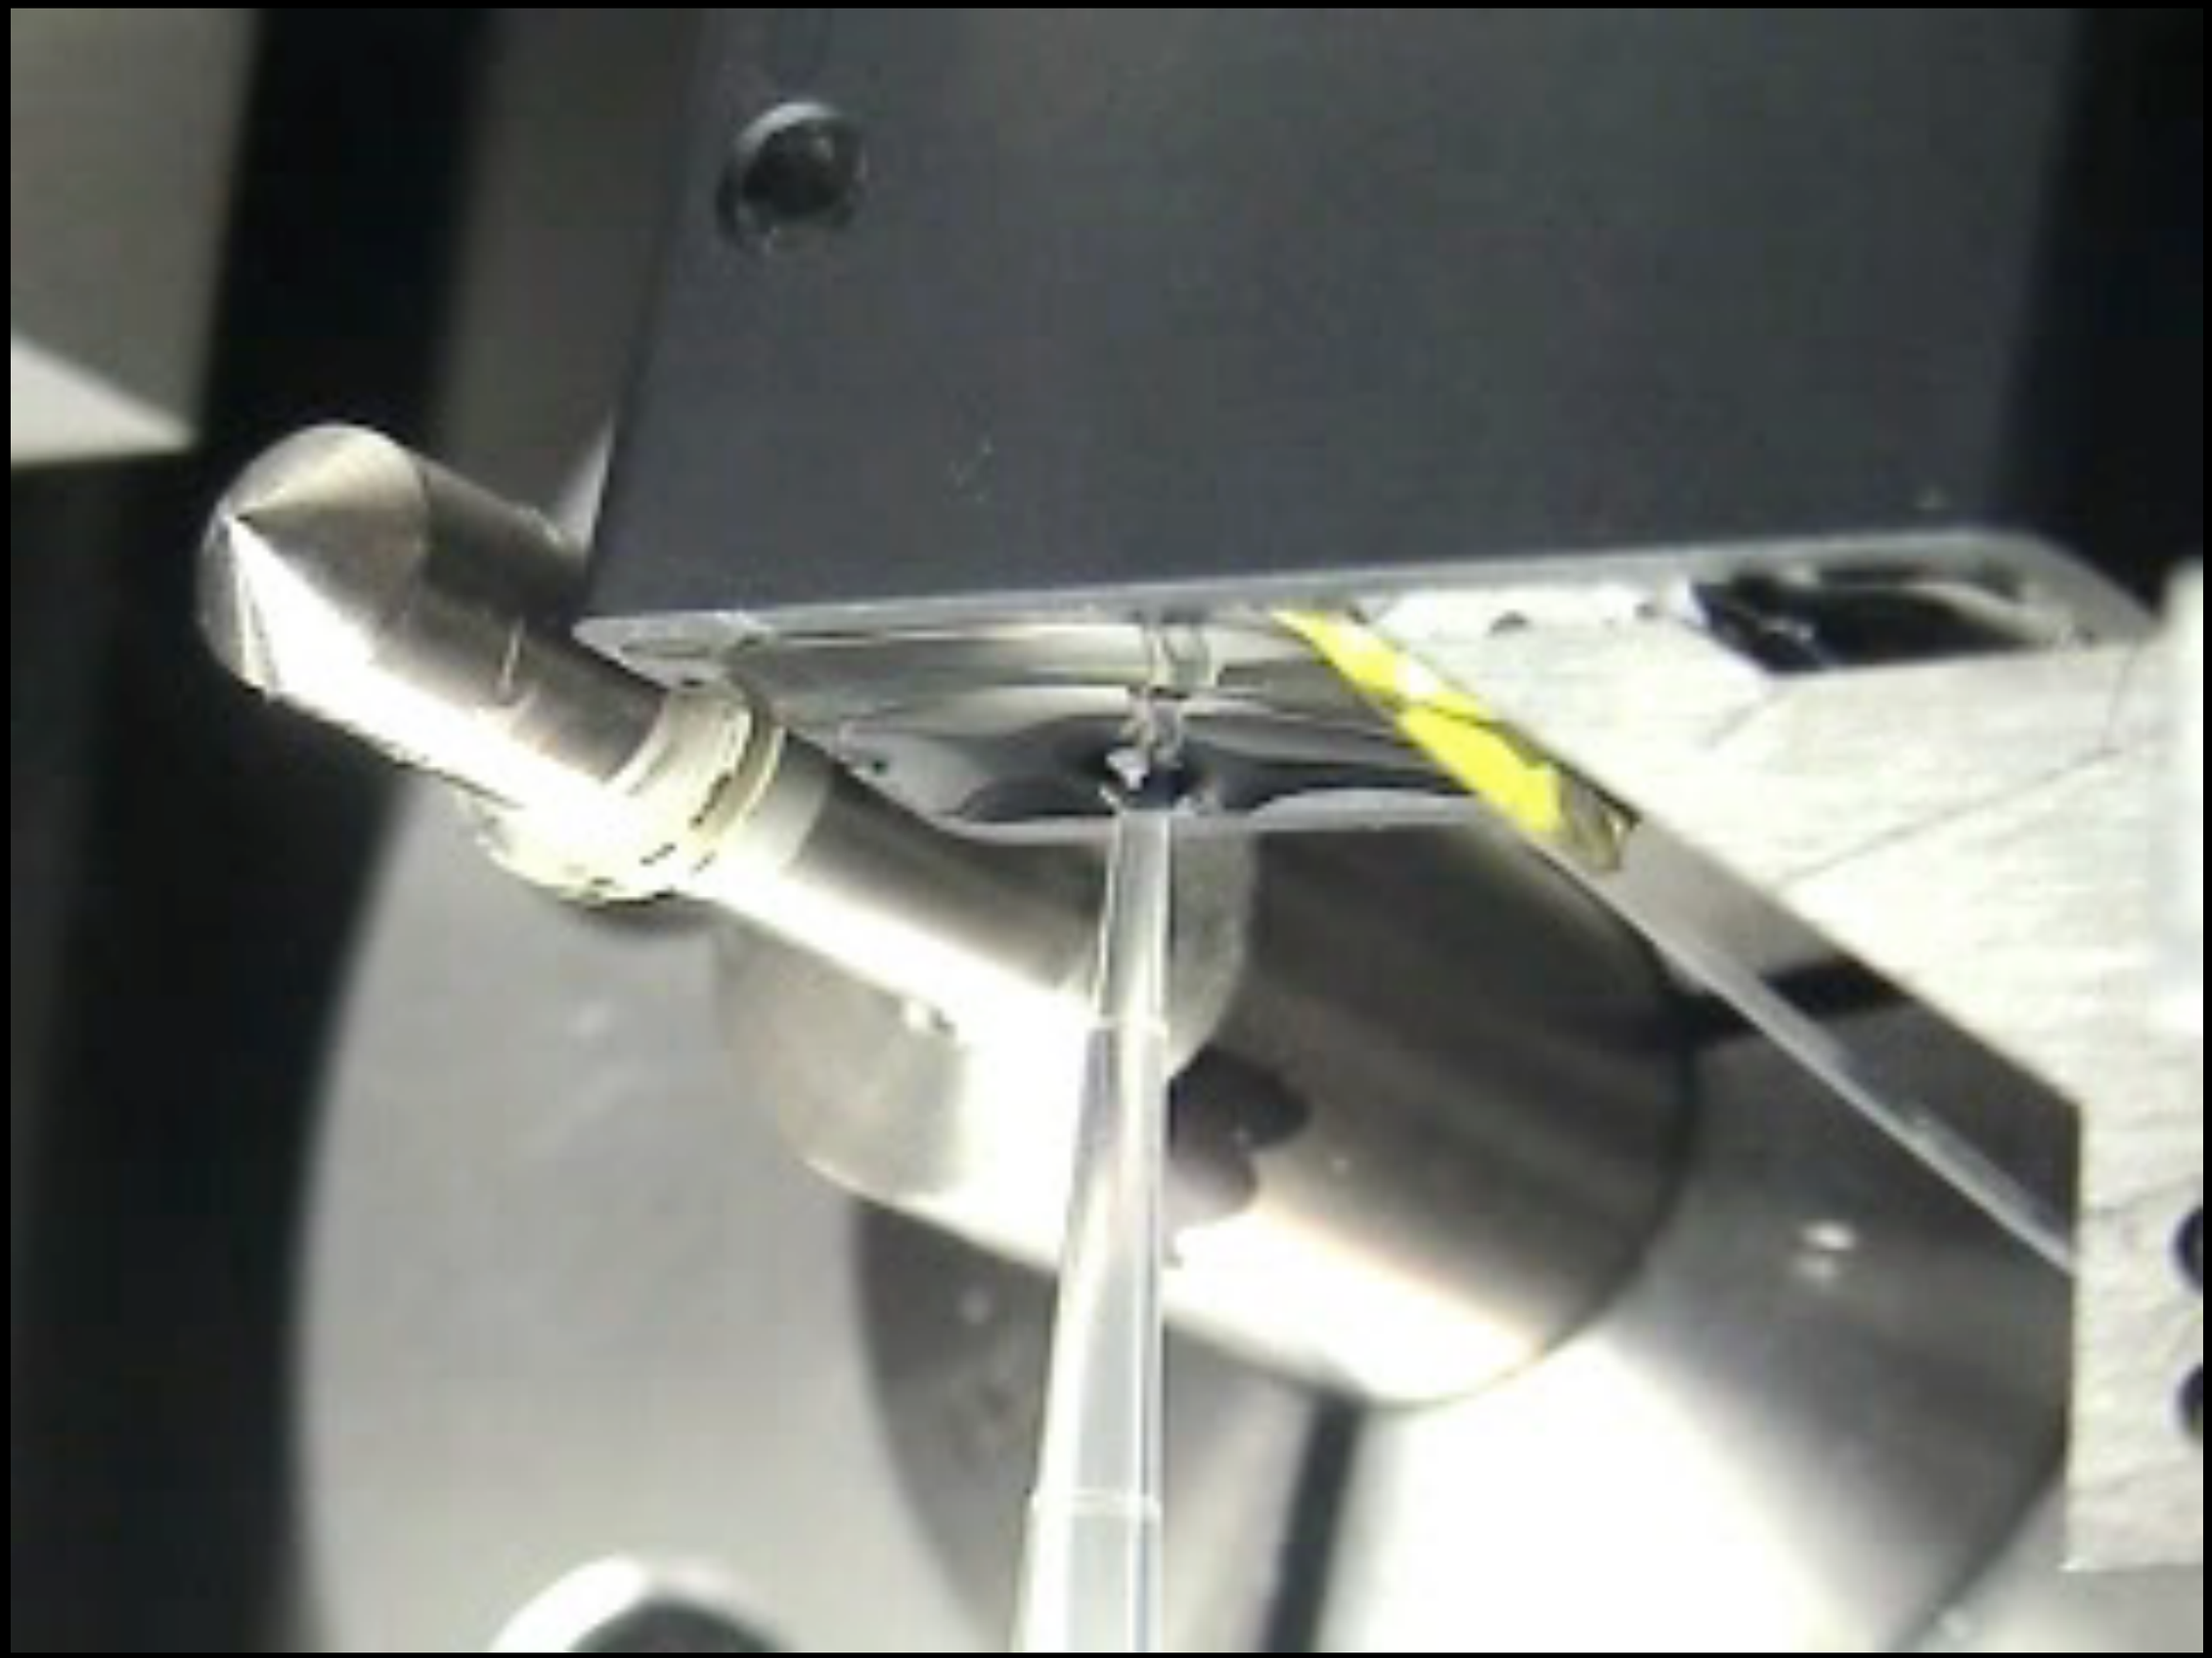

Inspiration for ATLUM

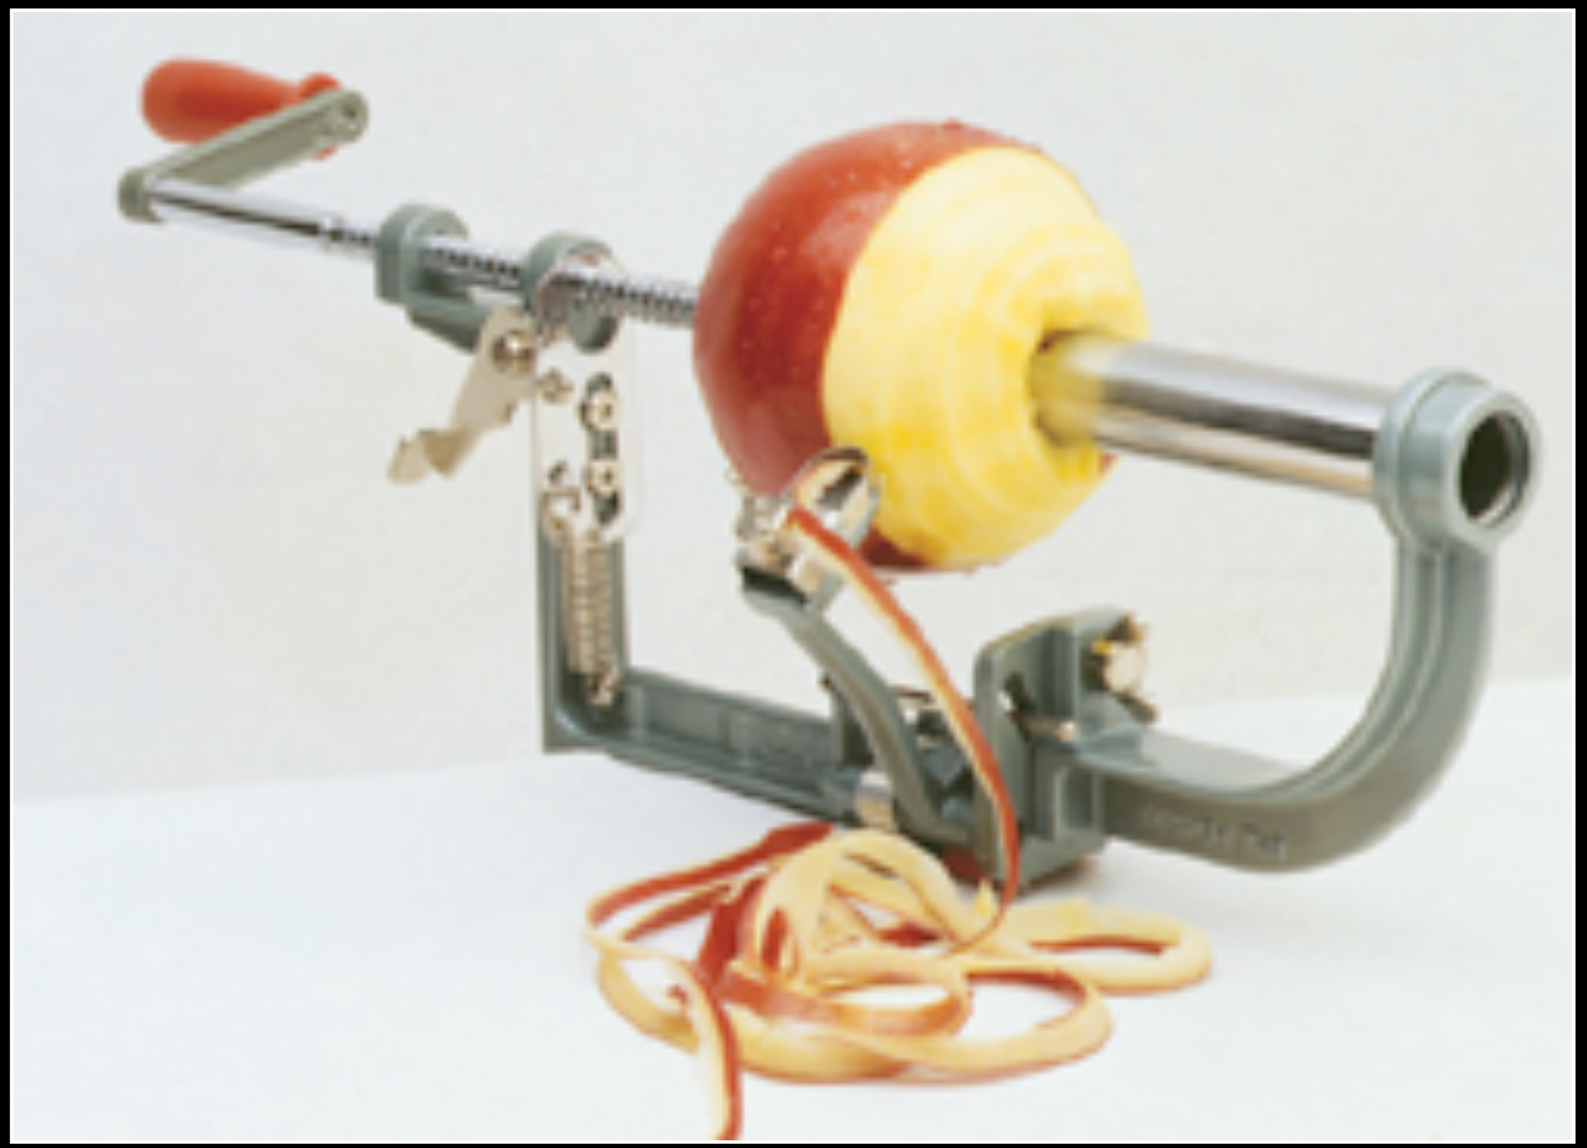

Ultrathin Section Libraries

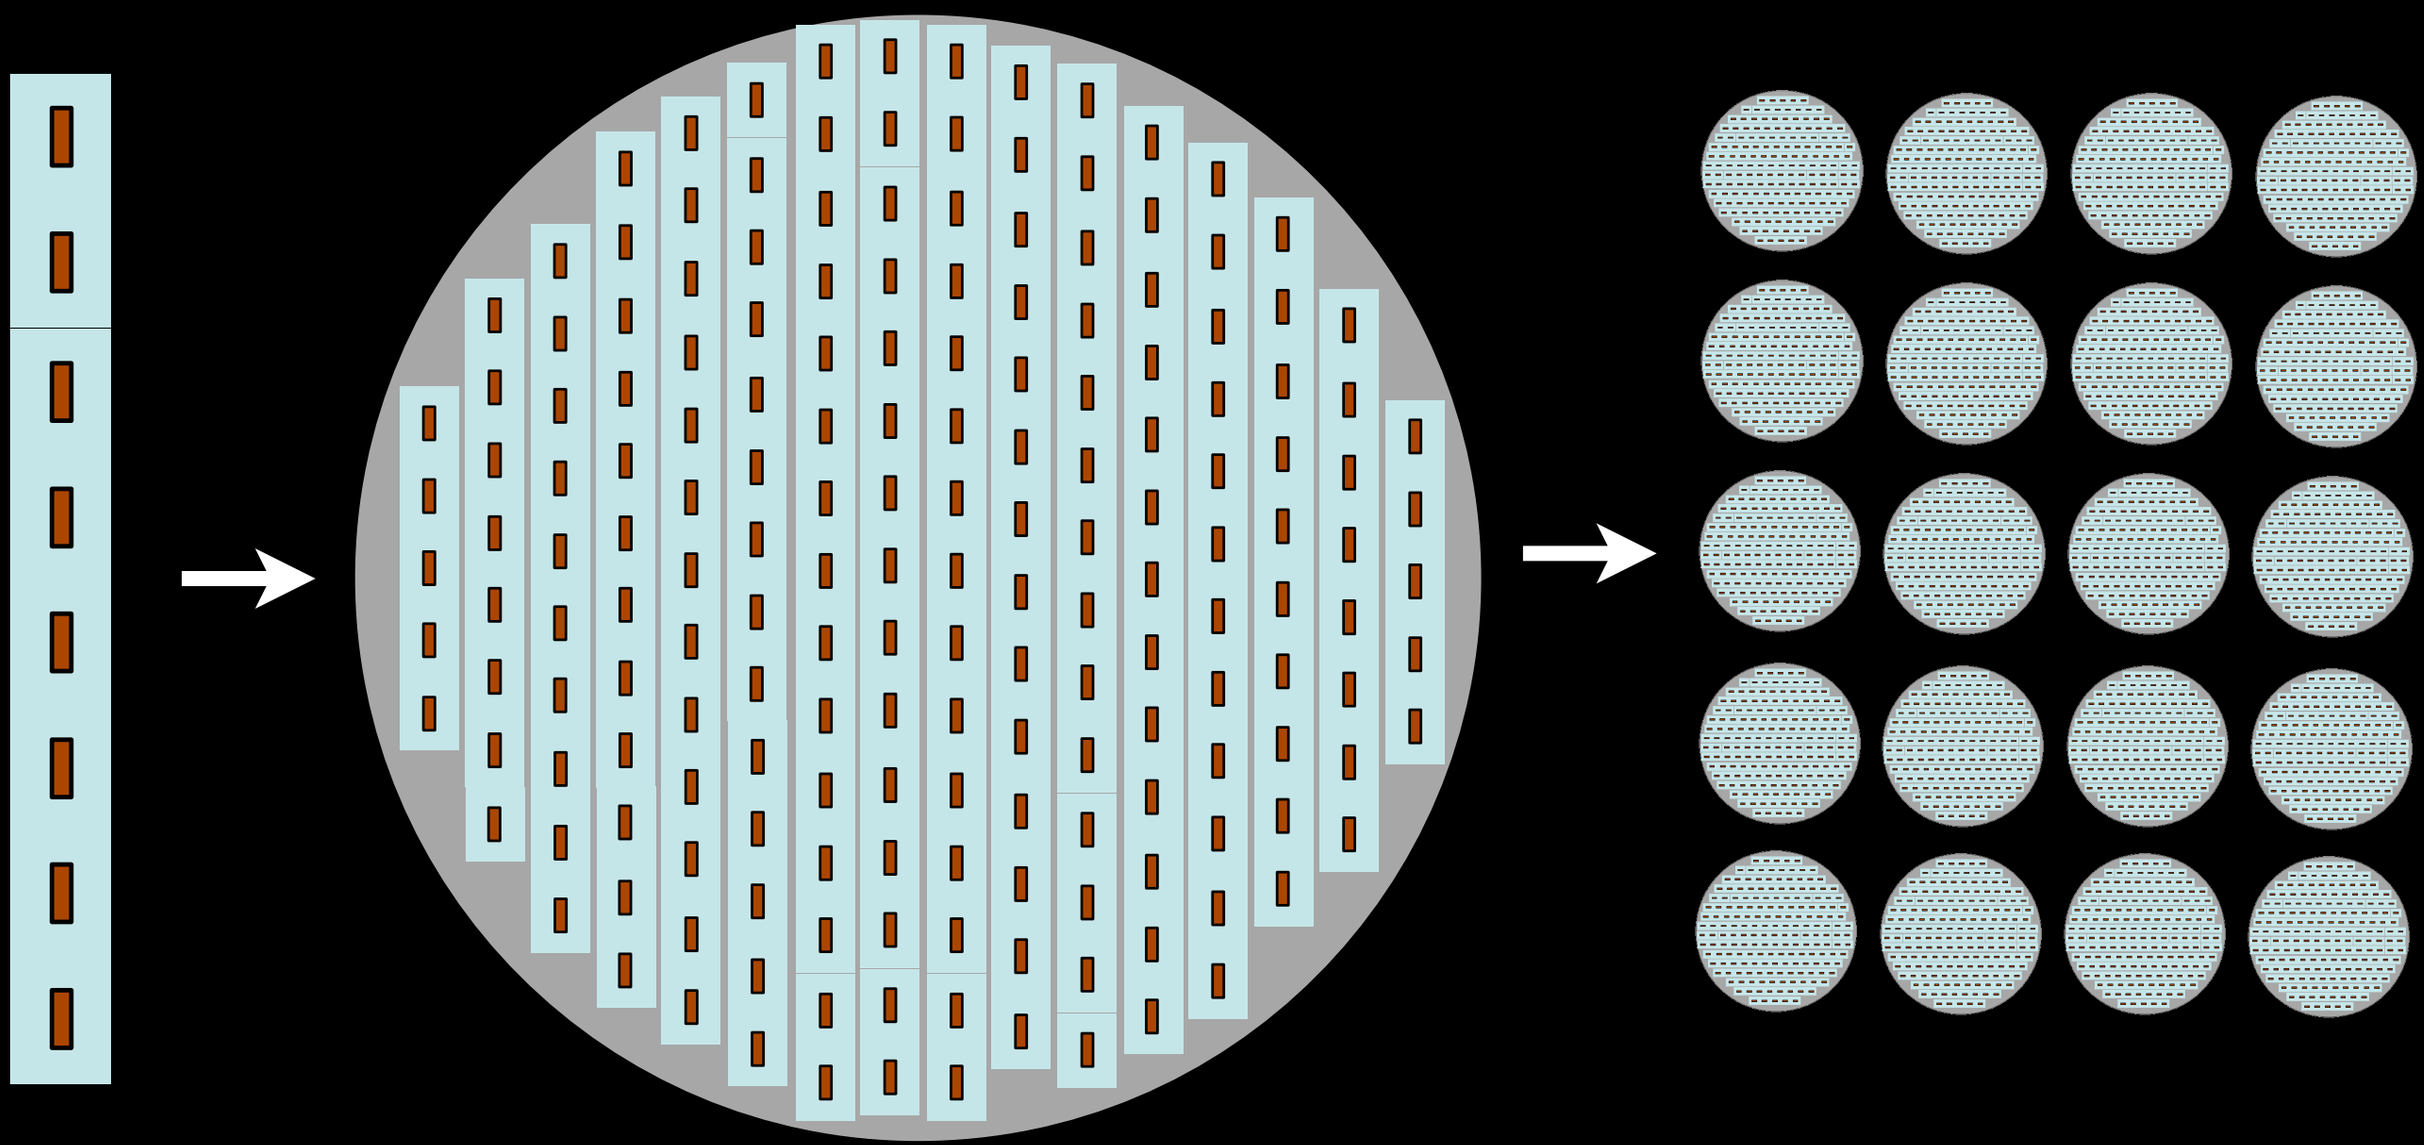

Connectome Pipeline

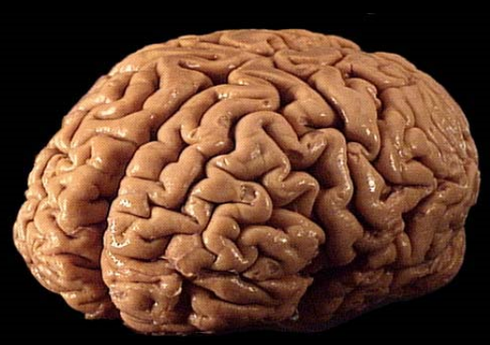

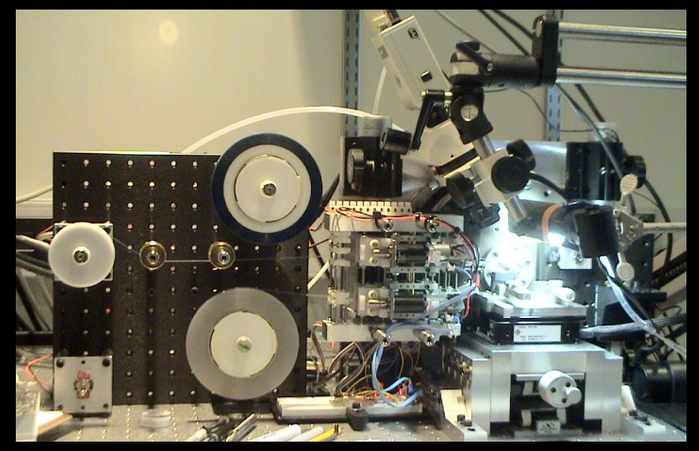

Get a brain → Get a piece → Cut it thin

Connectome Pipeline

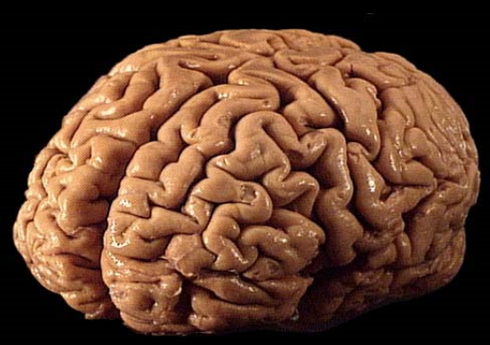

Get a brain

Get a piece

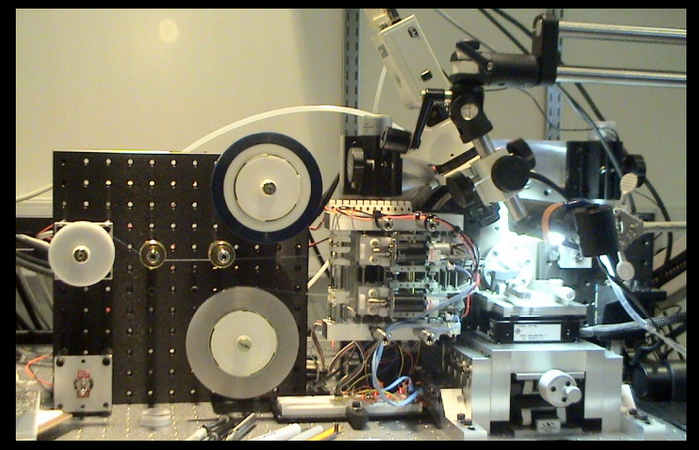

Cut it thin

Image it

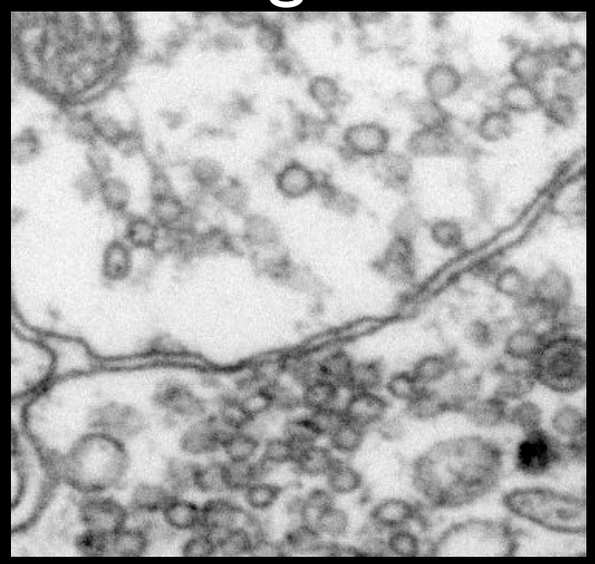

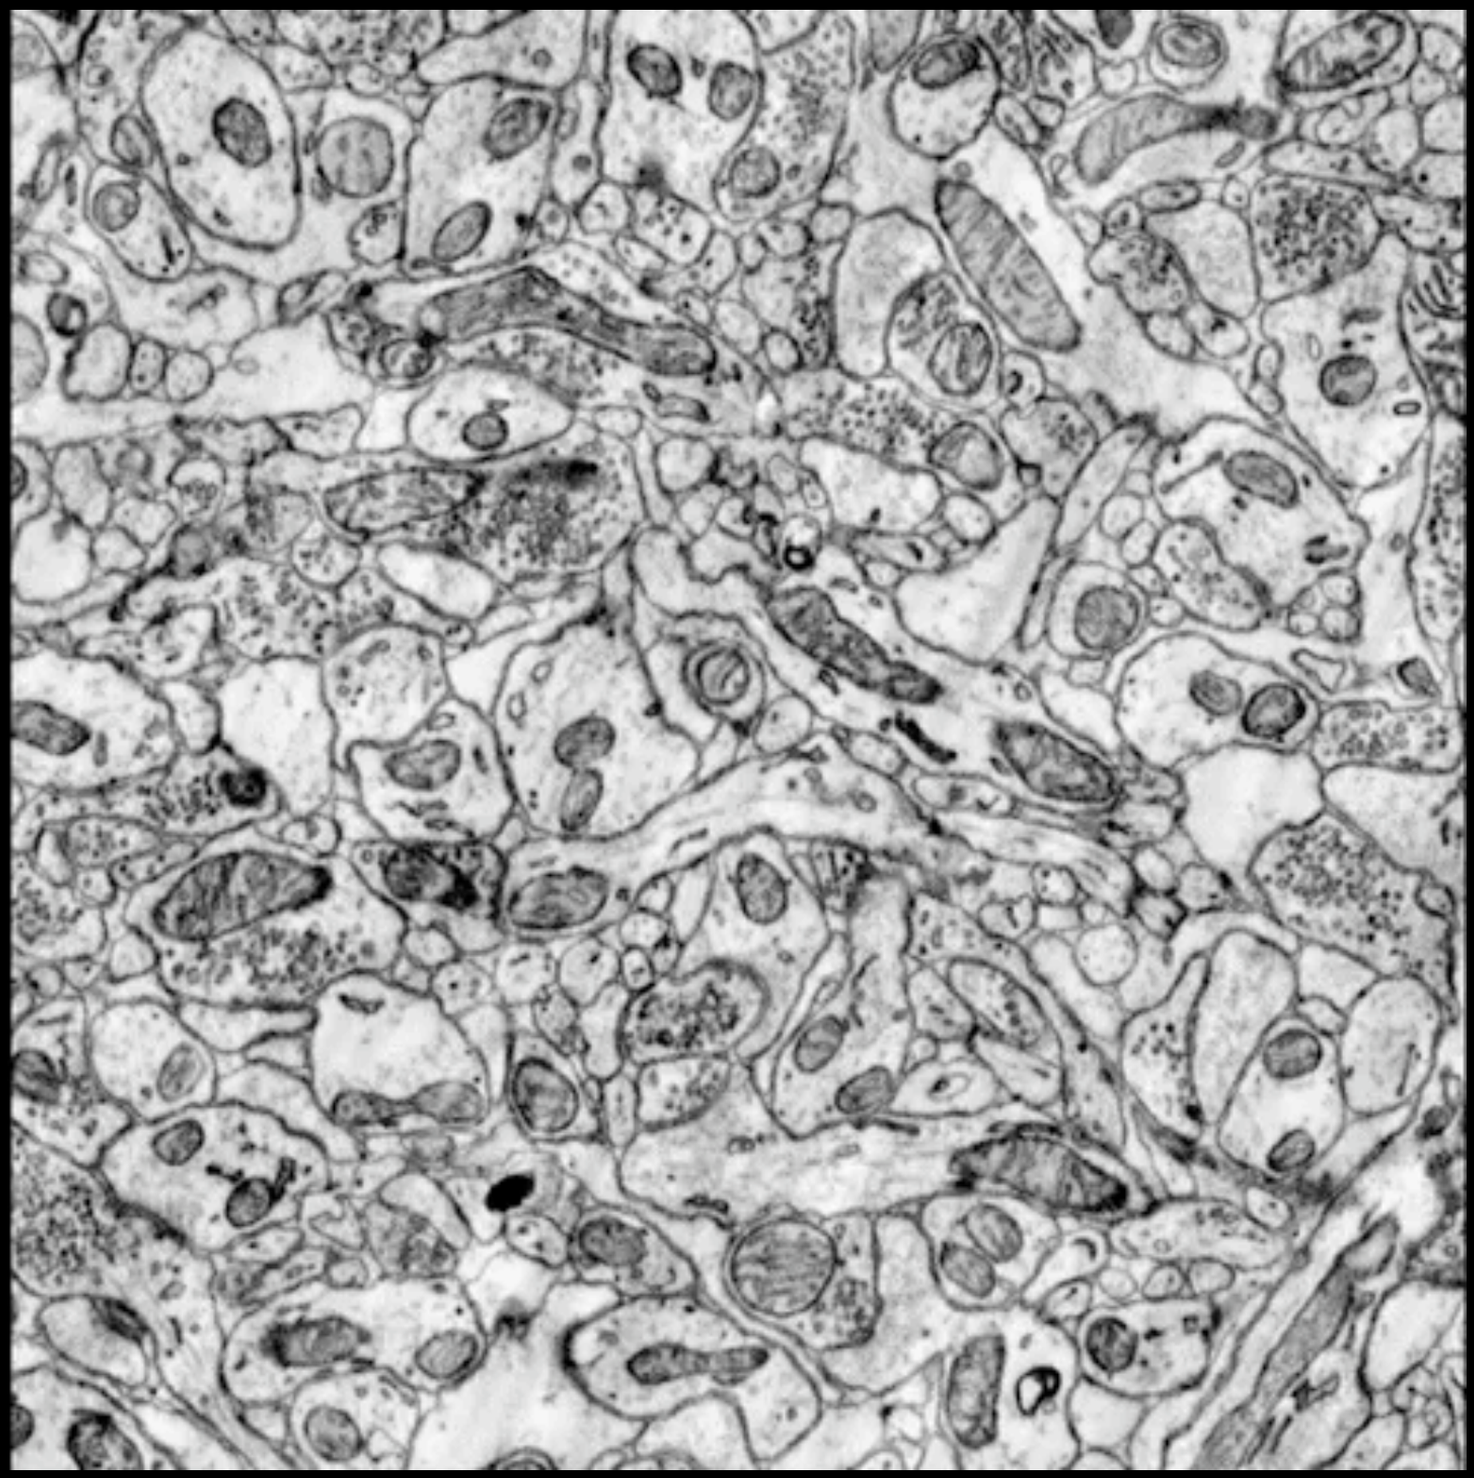

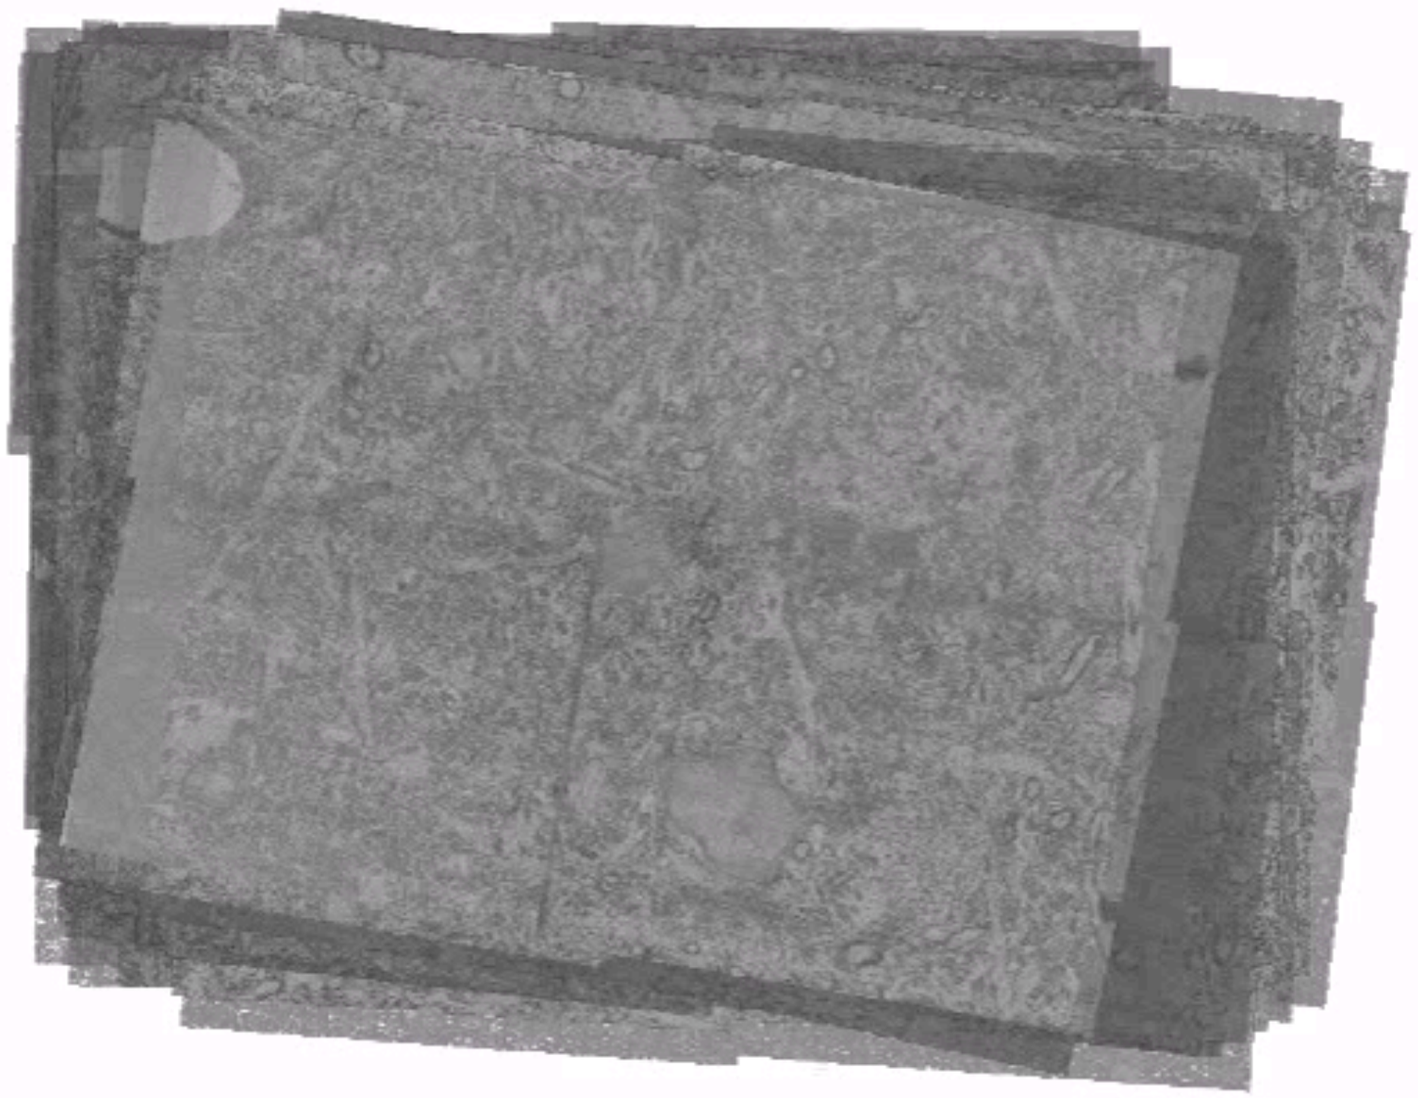

14,000 x 11,000 pixels, 300 sections = 40 GB

Imaging at 5 nm x 5 nm x 40 nm

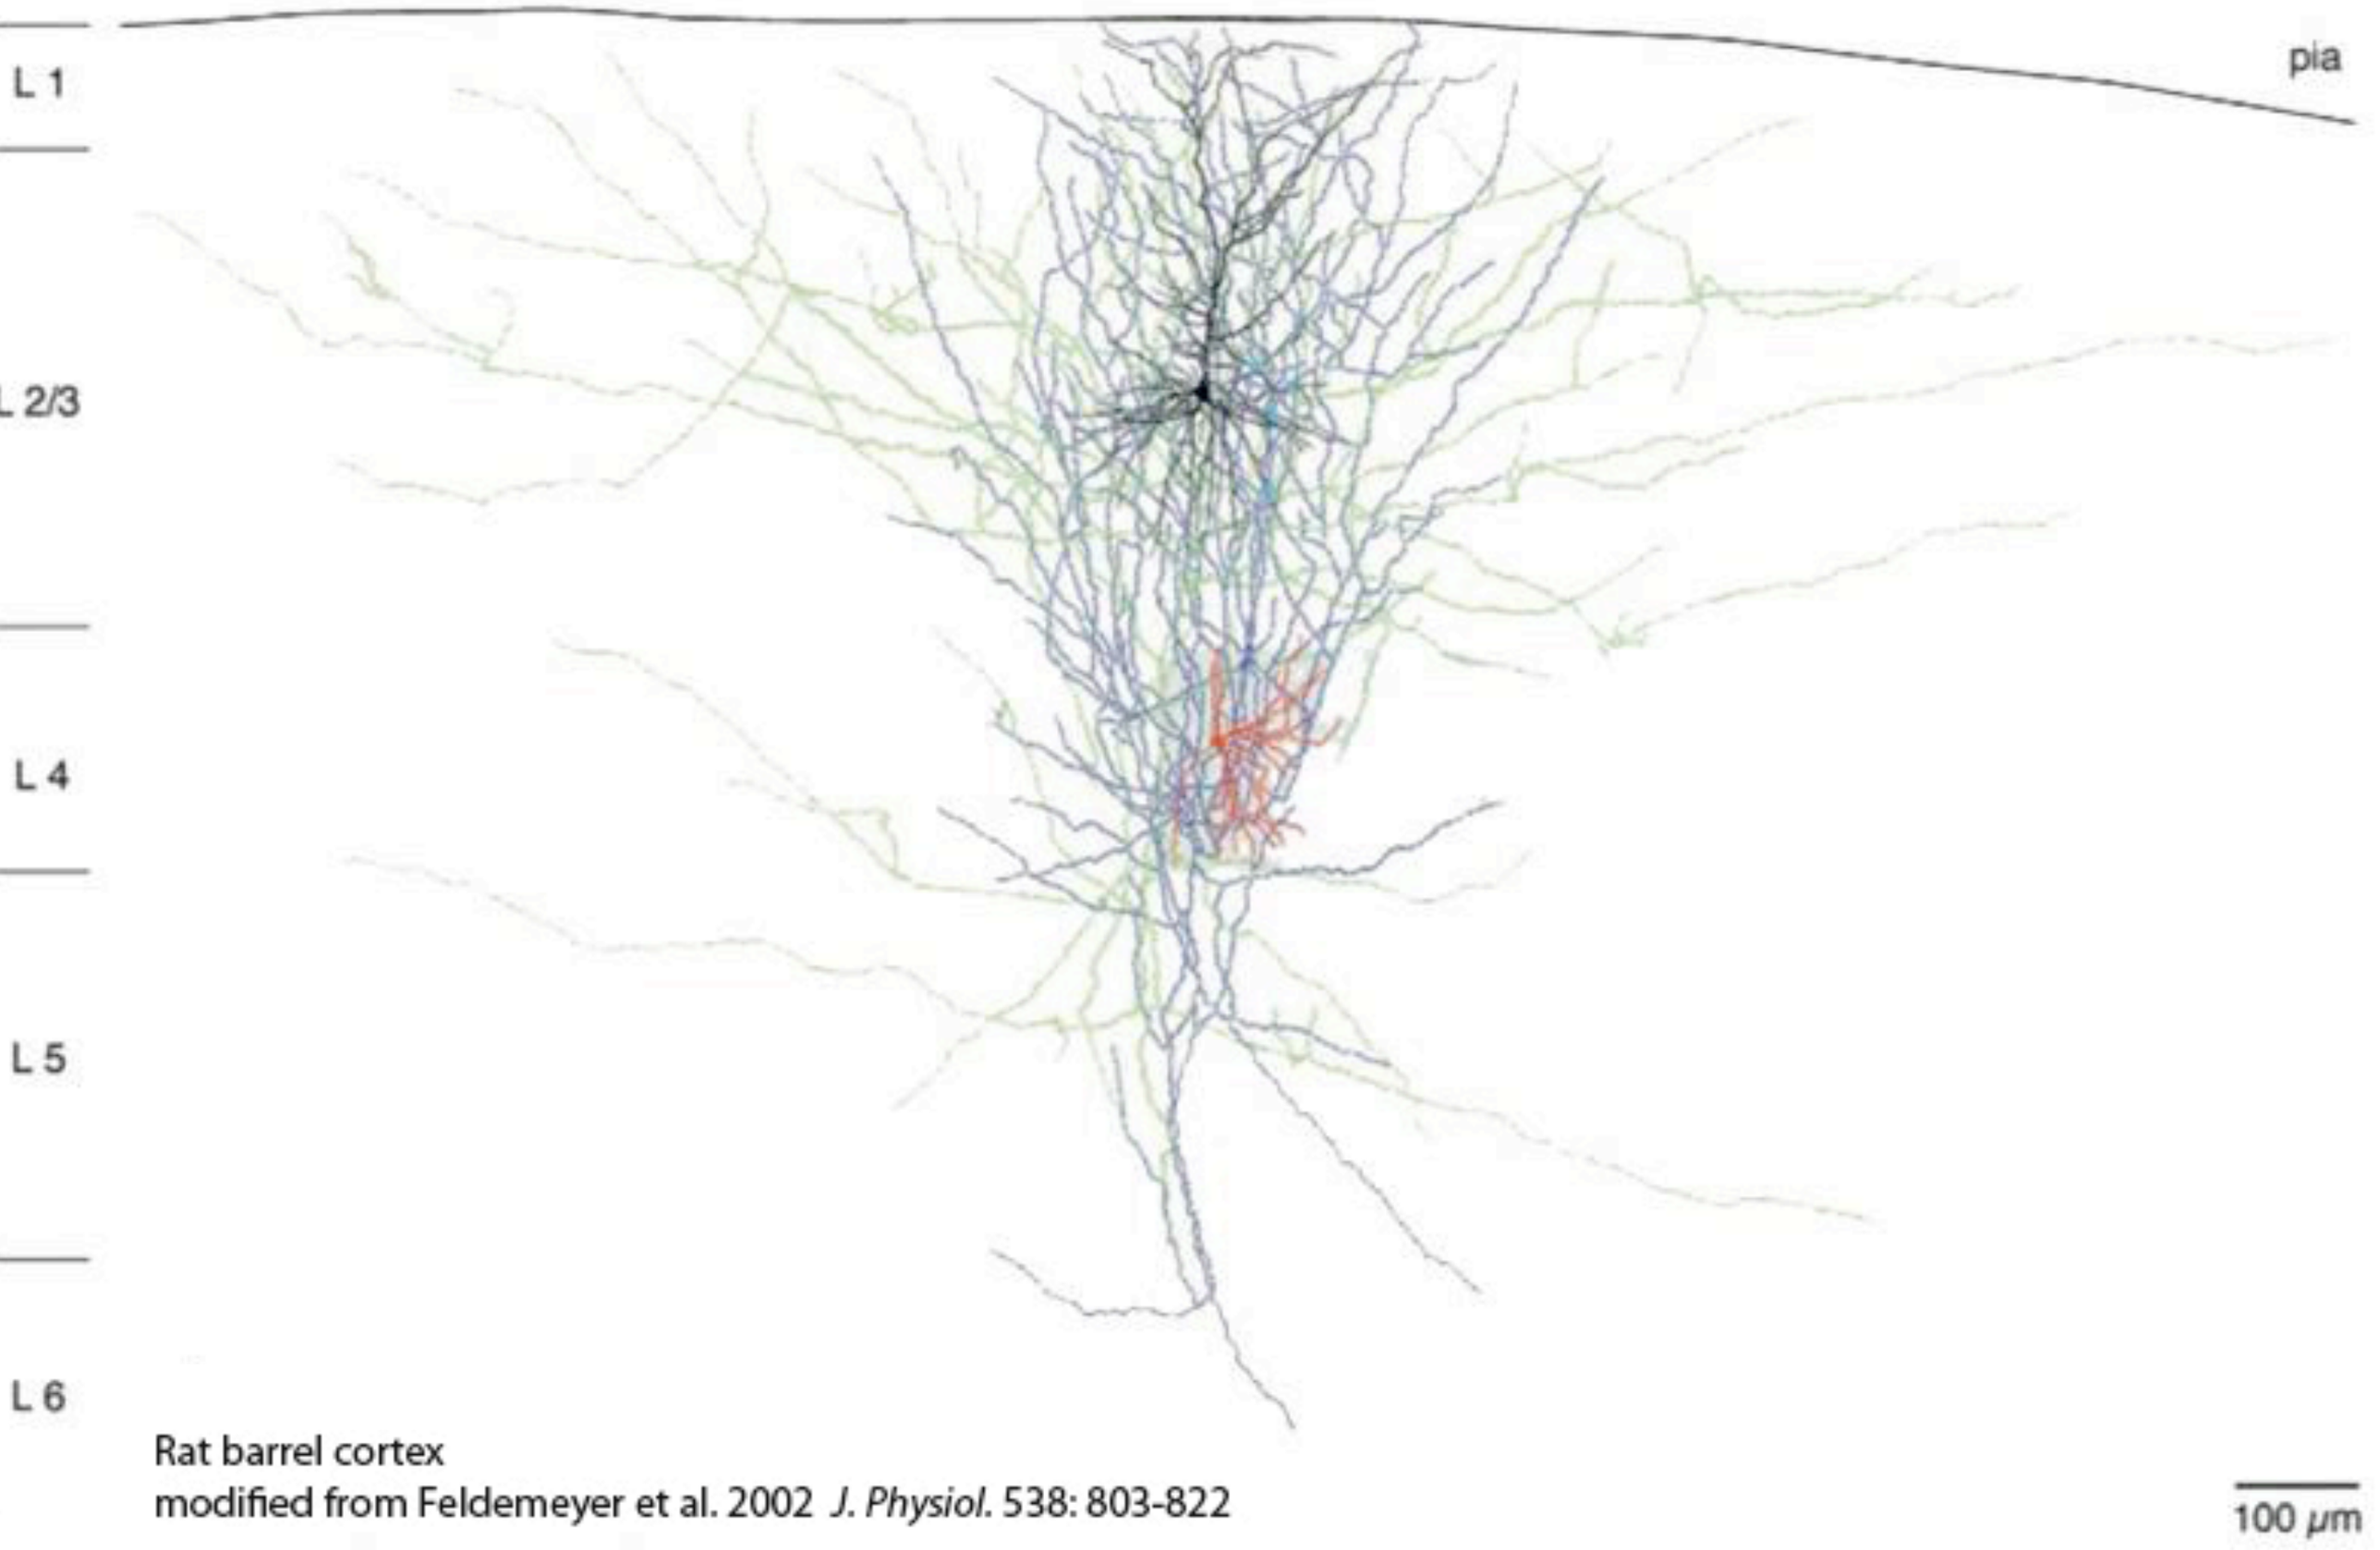

Rat barrel cortex
modified from Feldemeyer et al. 2002 *J. Physiol.* 538: 803-822

Imaging at 5 nm x 5 nm x 40 nm

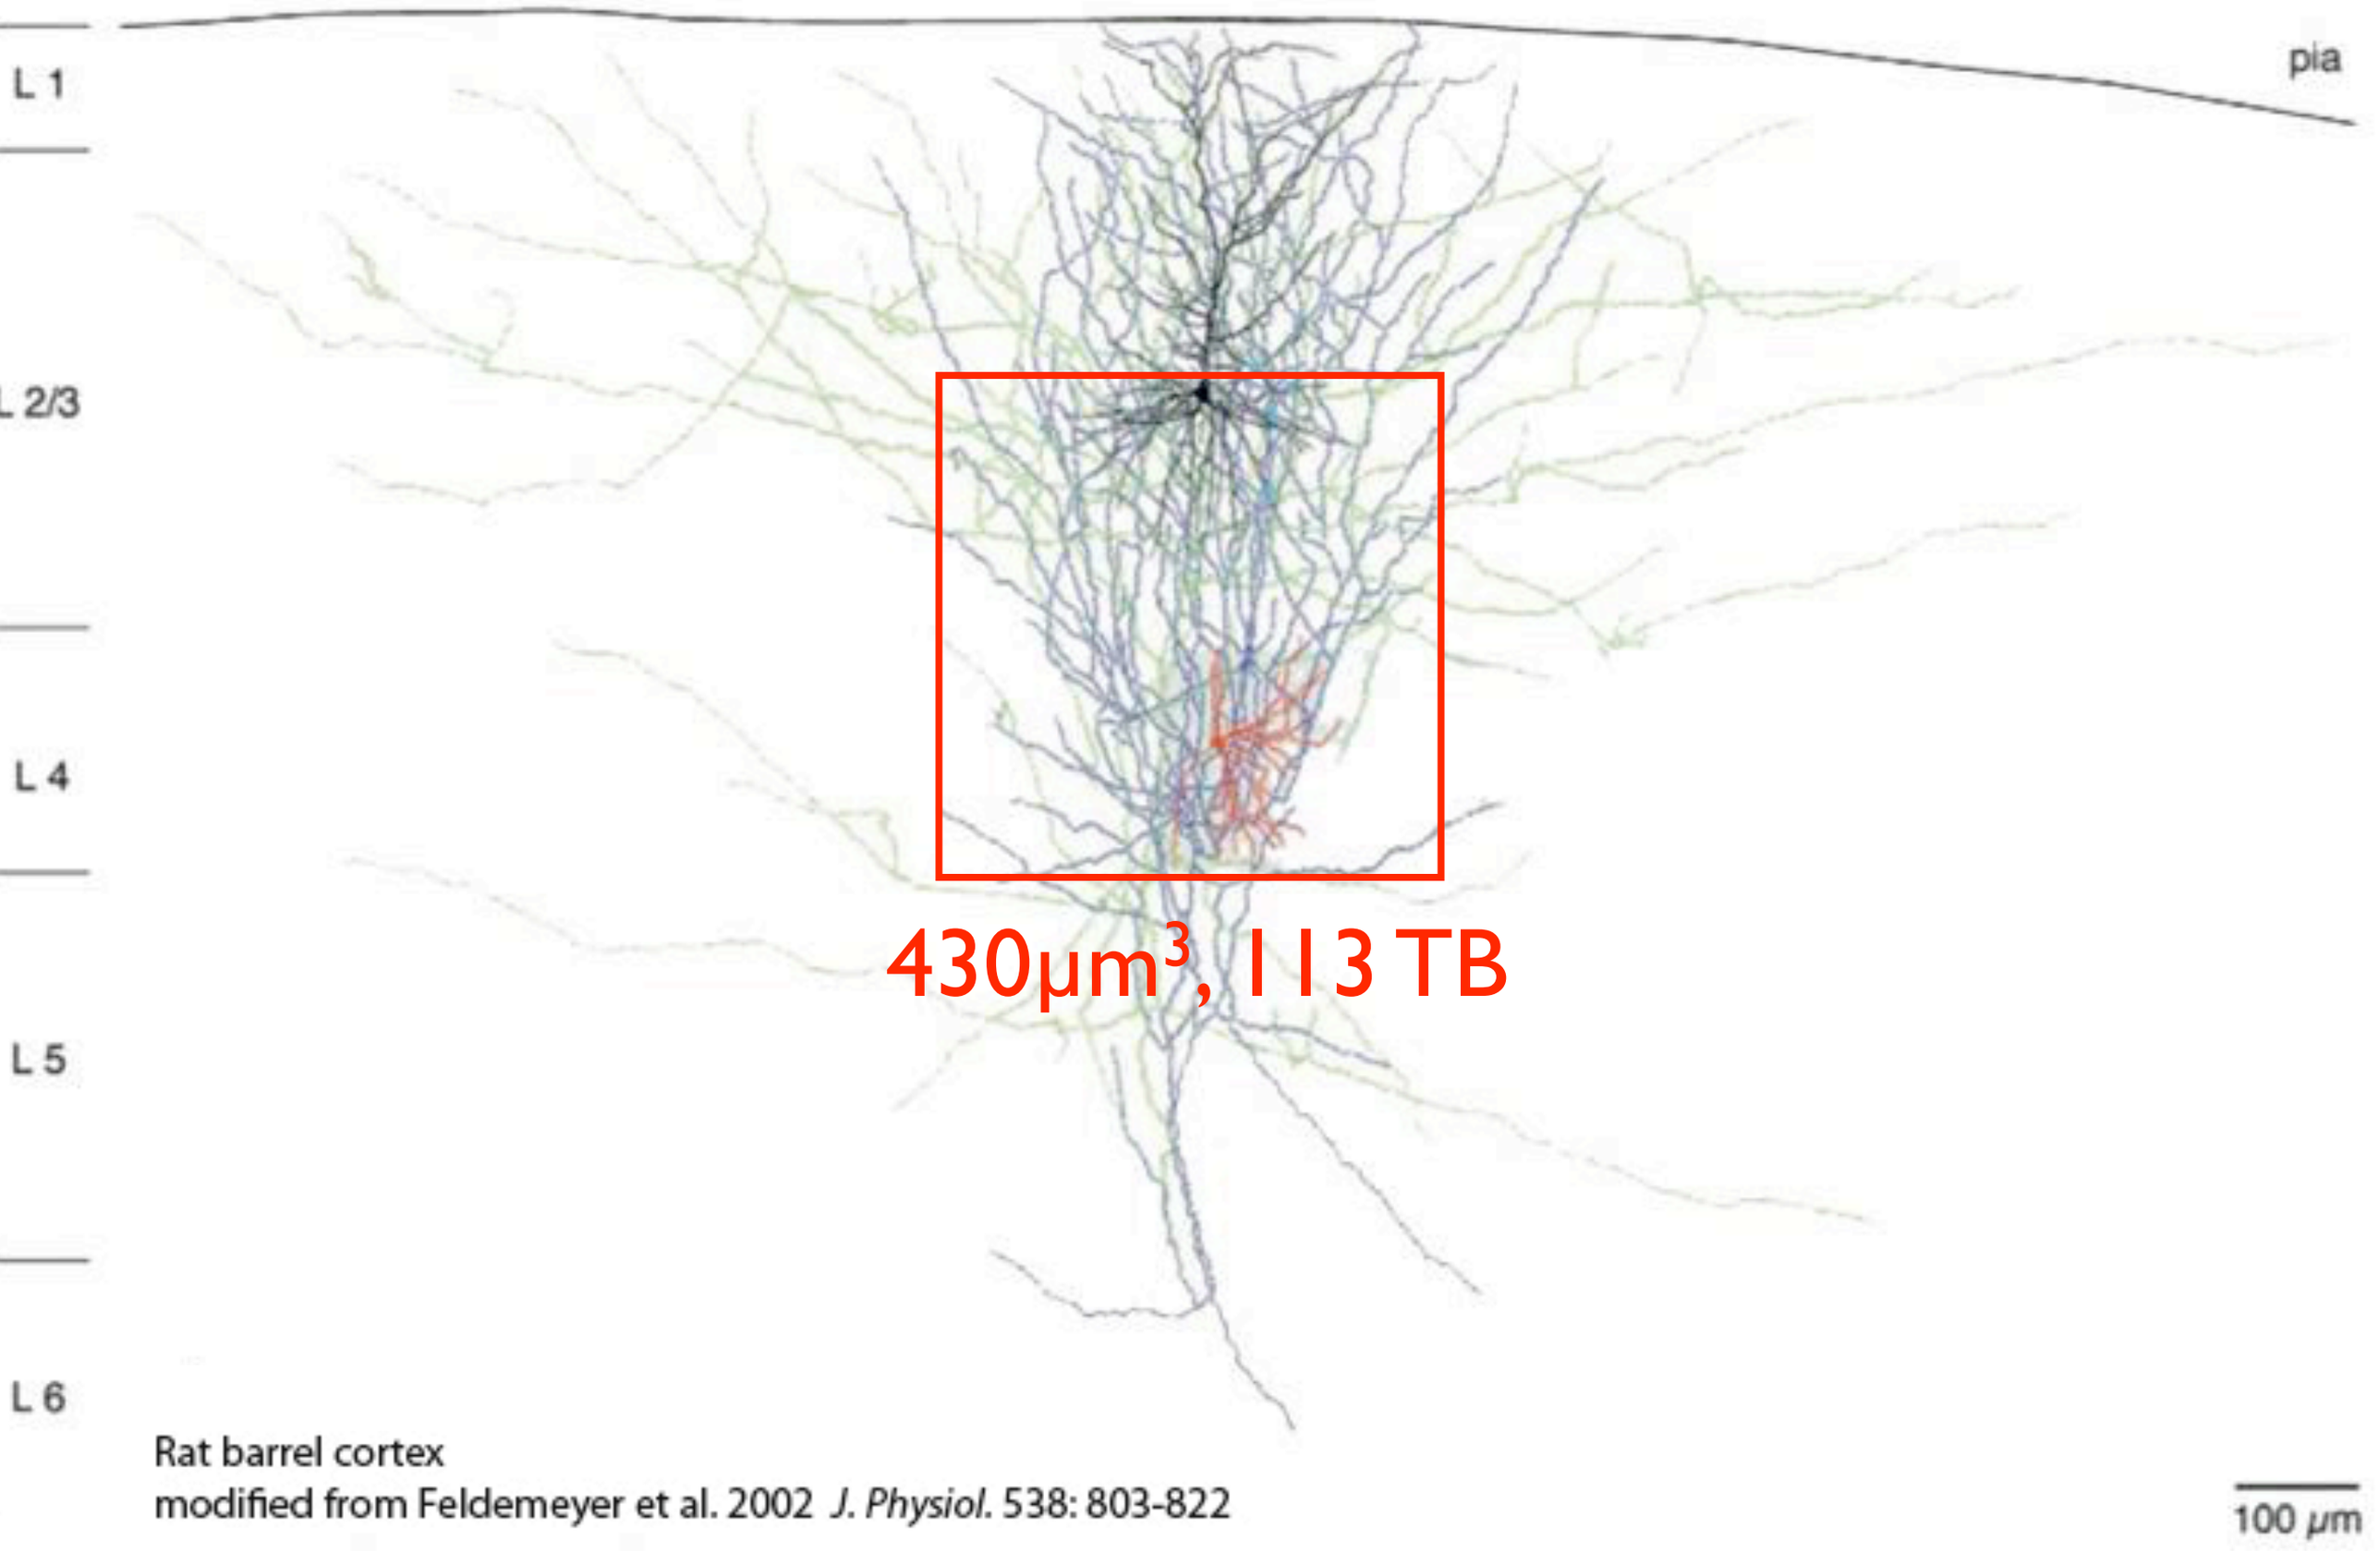

Imaging at 5 nm x 5 nm x 40 nm

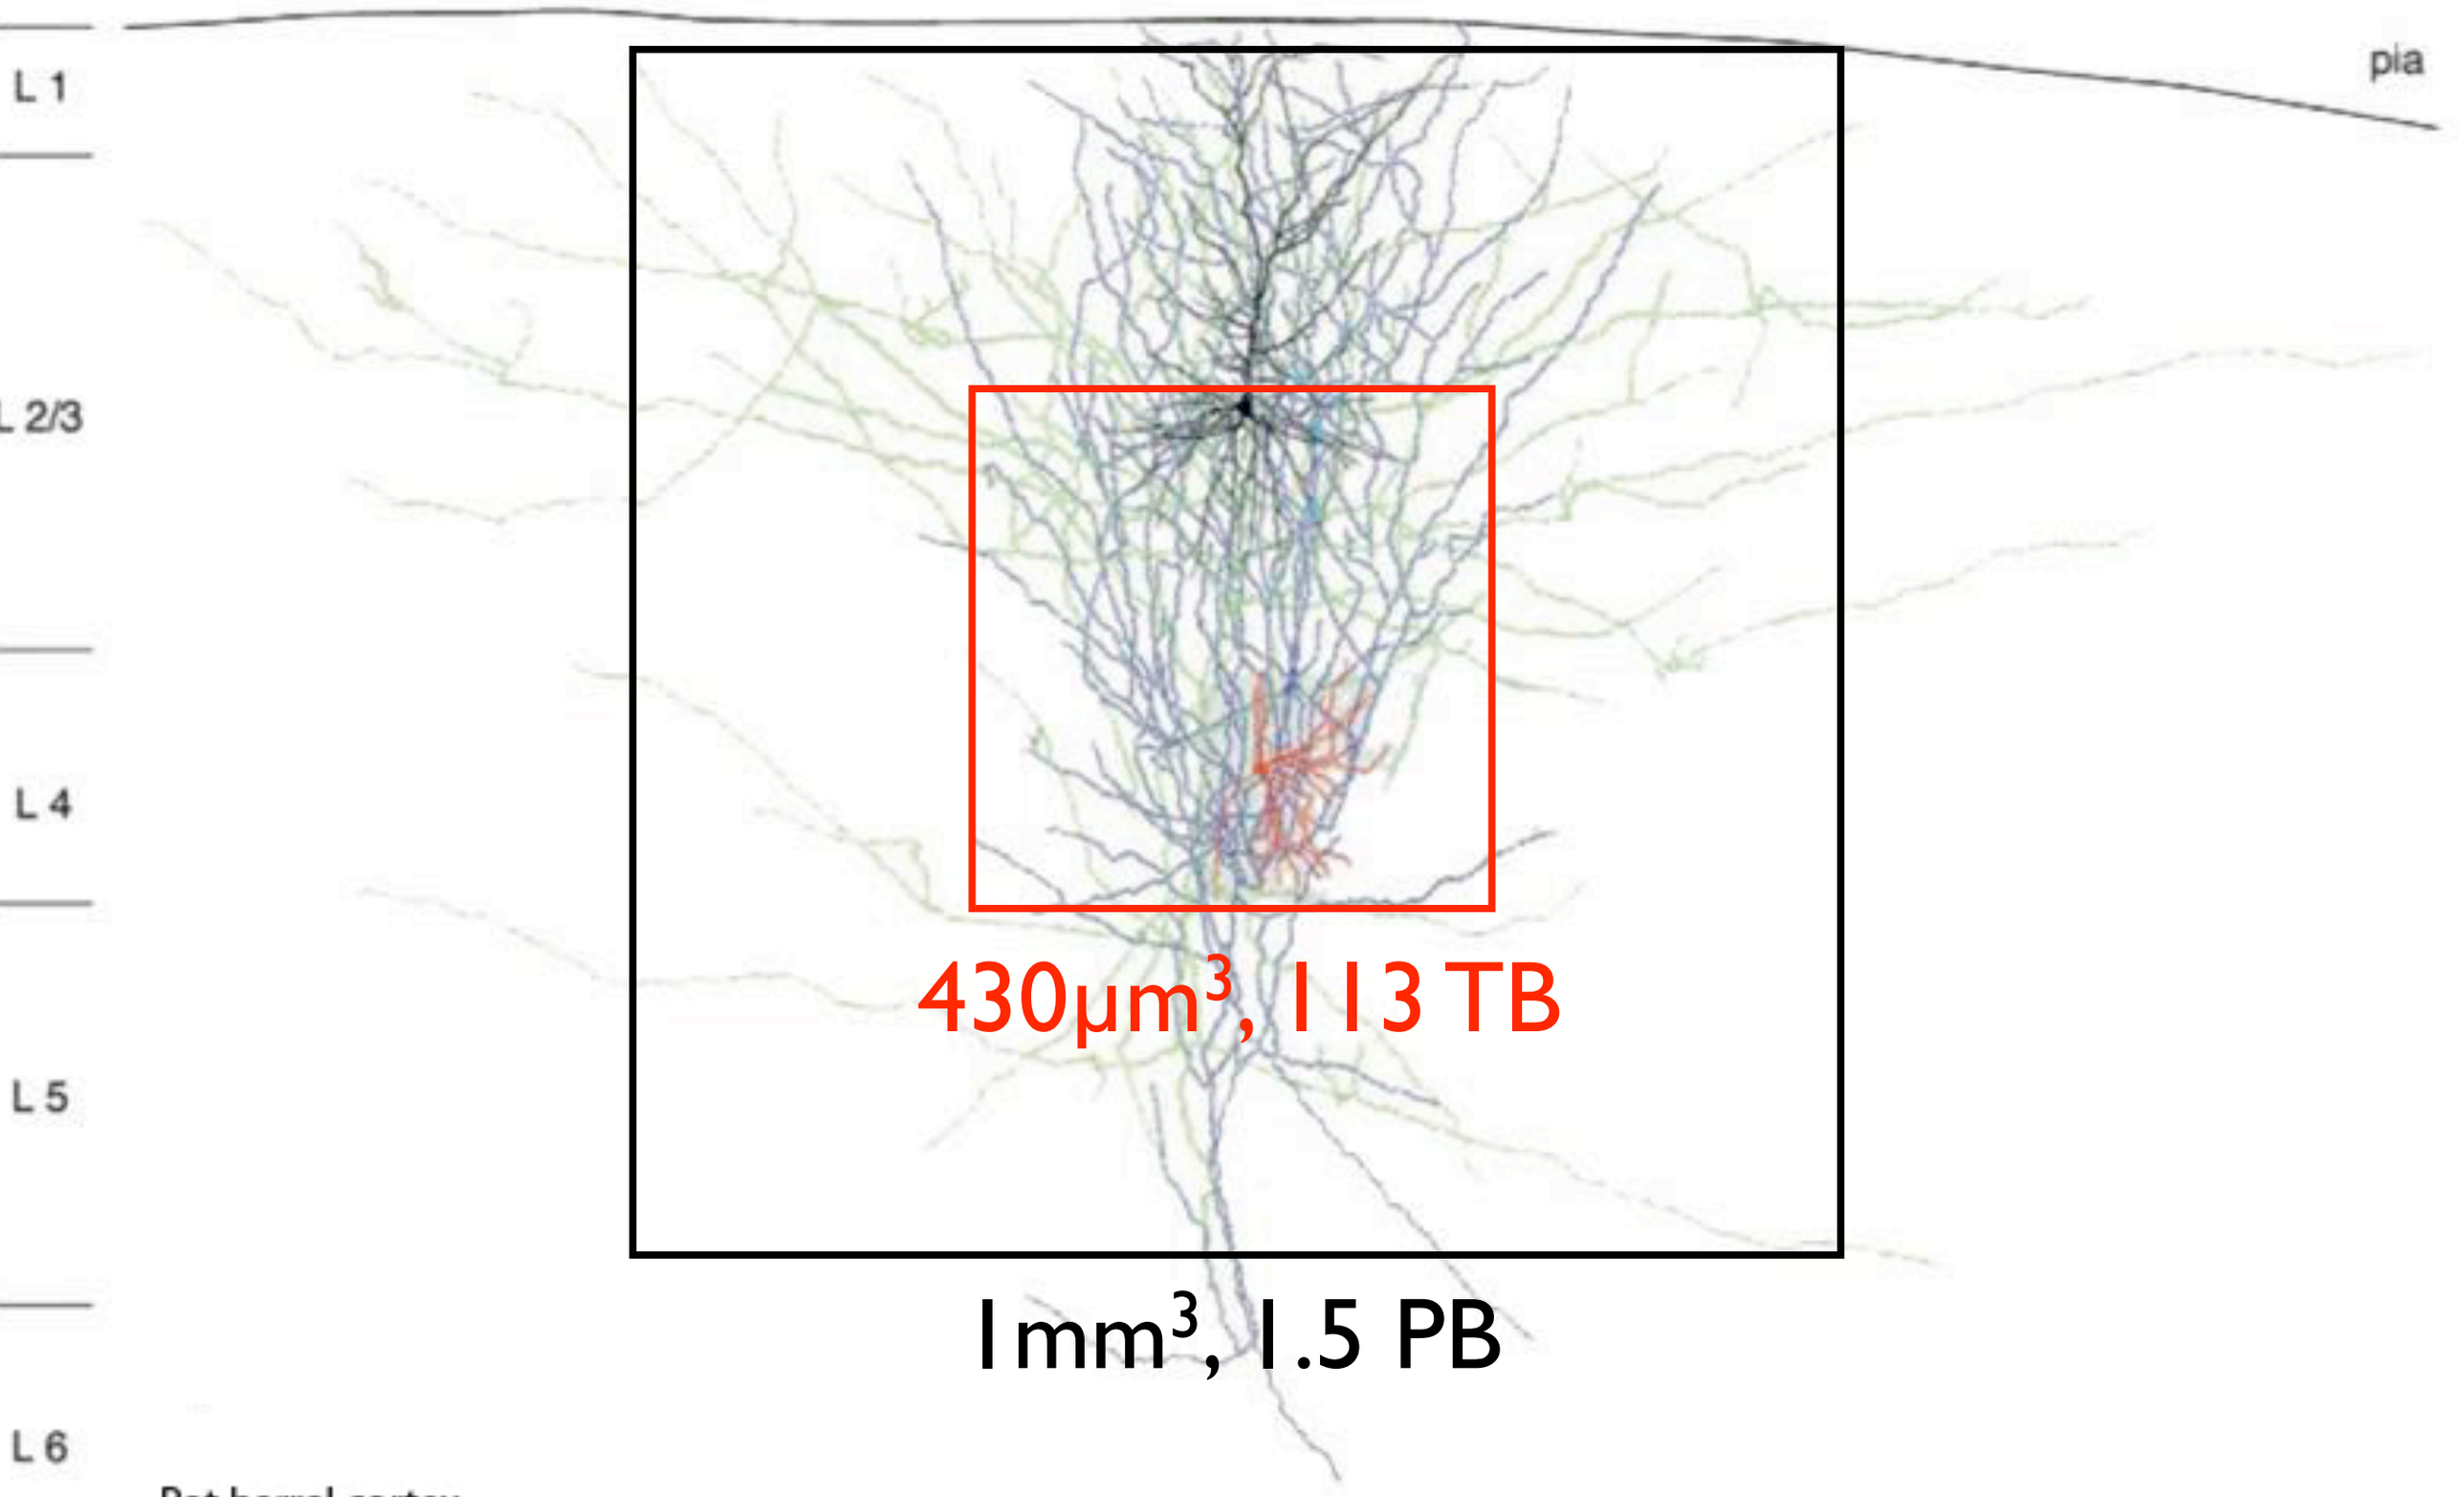

Connectome Pipeline

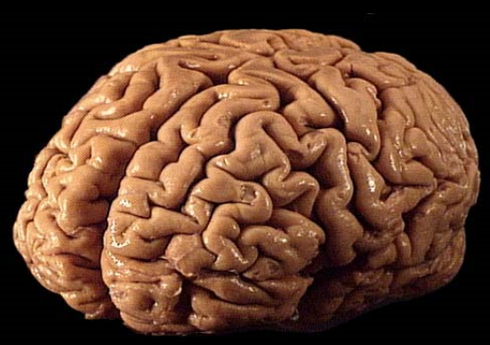

Get a brain

Get a piece

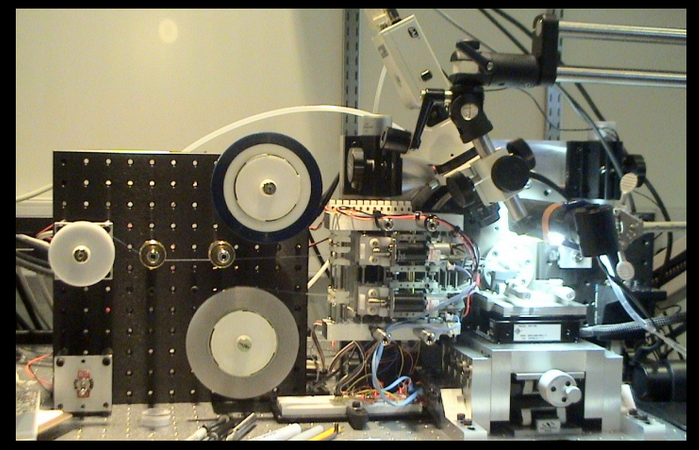

Cut it thin

Image it

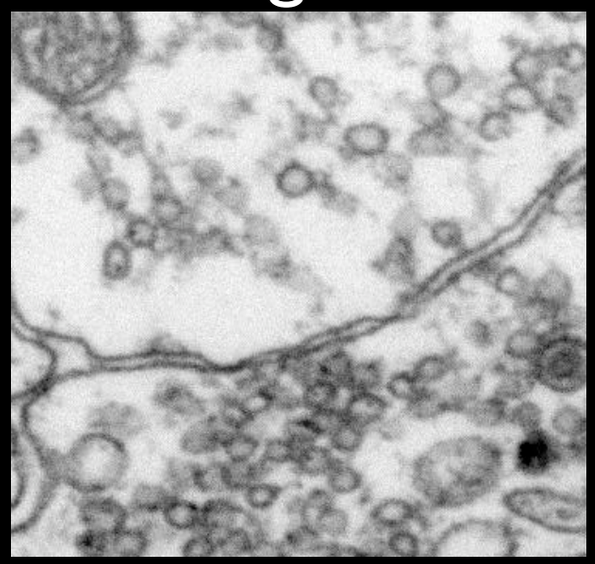

Connectome Pipeline

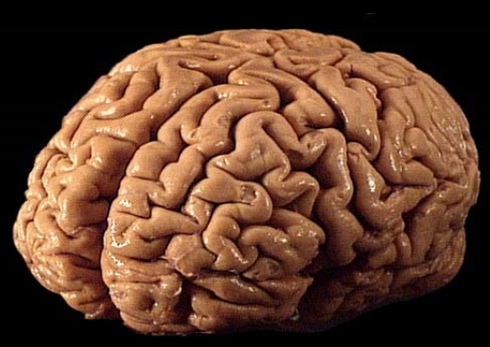

Get a brain

Get a piece

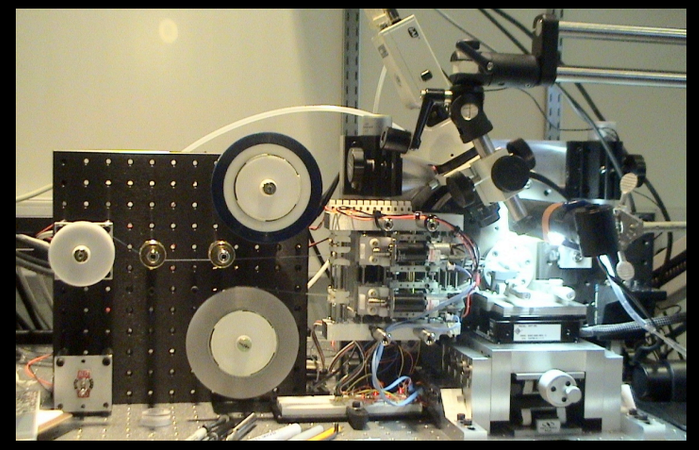

Cut it thin

Image it

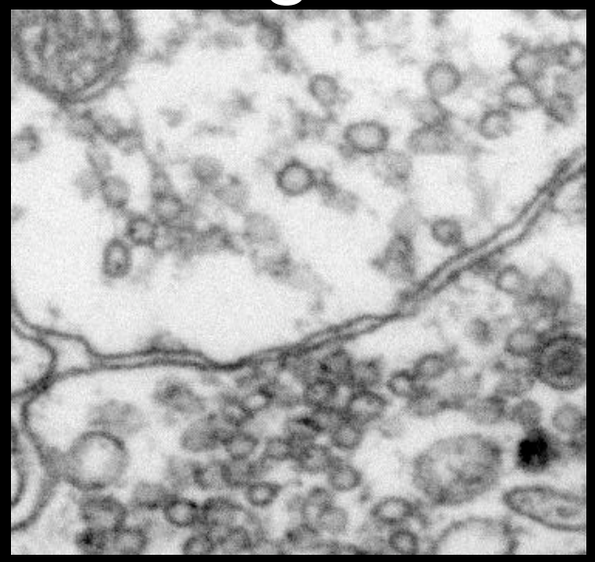

Reconstruct

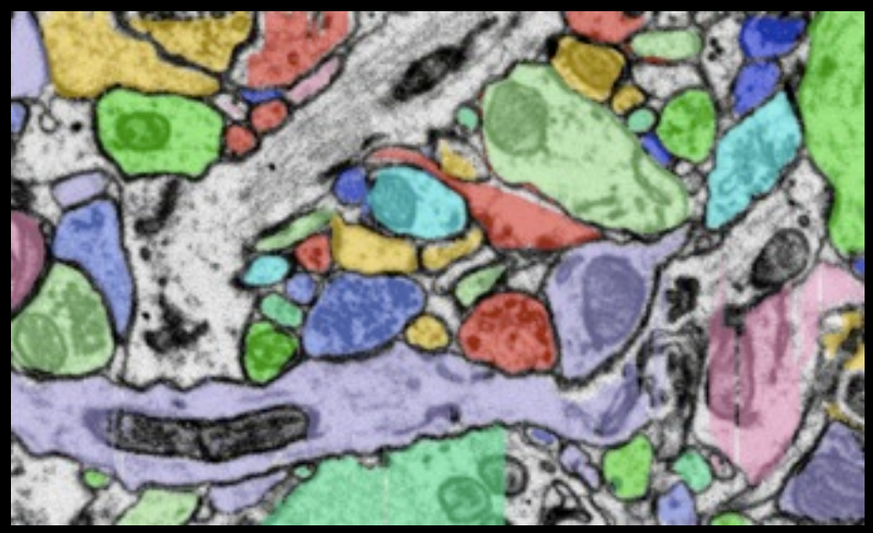

Connectome Pipeline

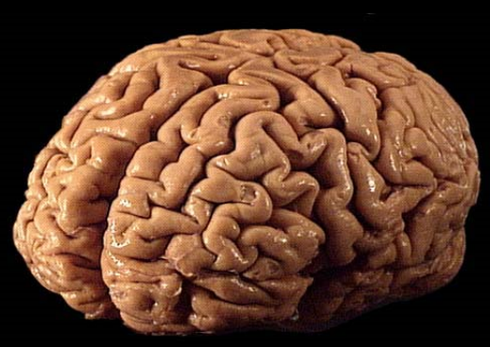

Get a brain →

Get a piece →

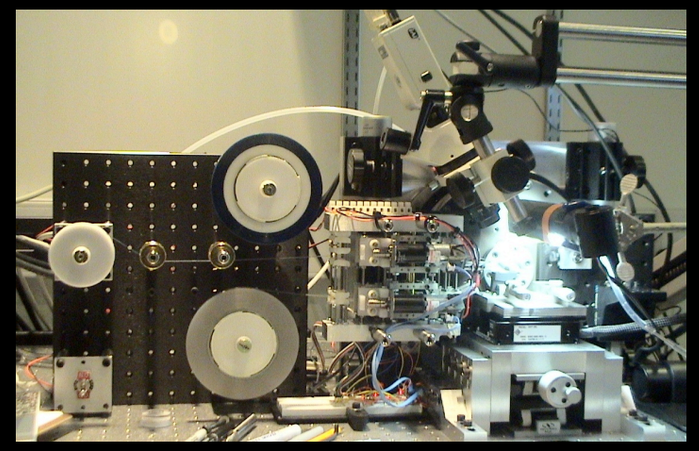

Cut it thin ↓

Image it ←

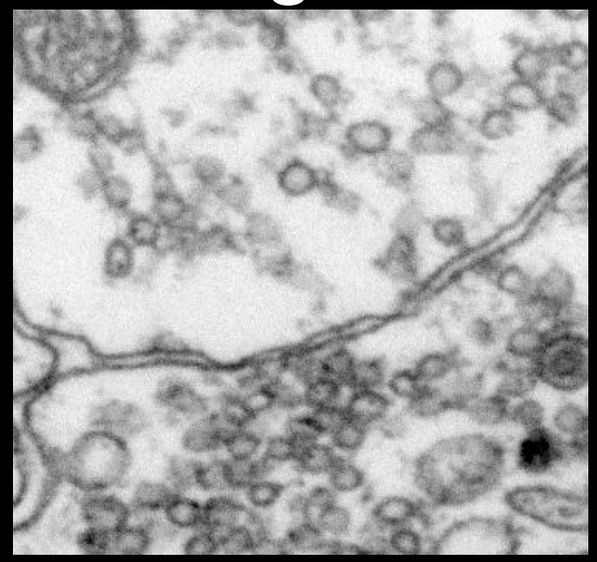

Reconstruct ←

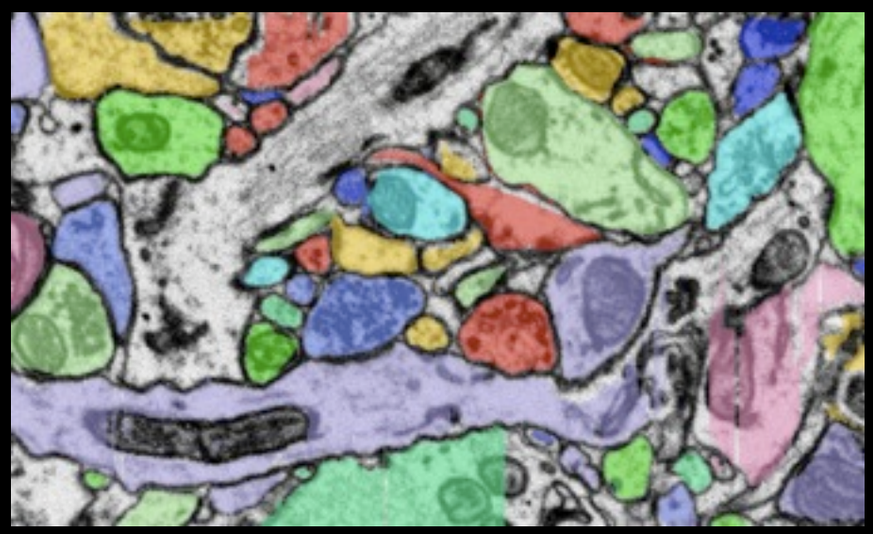

Visualize & Analyze ↔

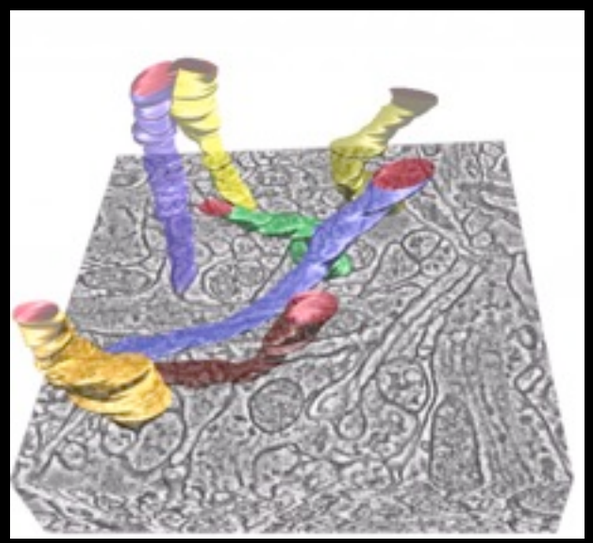

Reconstruction

- Active Ribbons [Vazquez 09]

Section 1

GPU Processing

- Active Ribbons on GPUs: 23x speedup

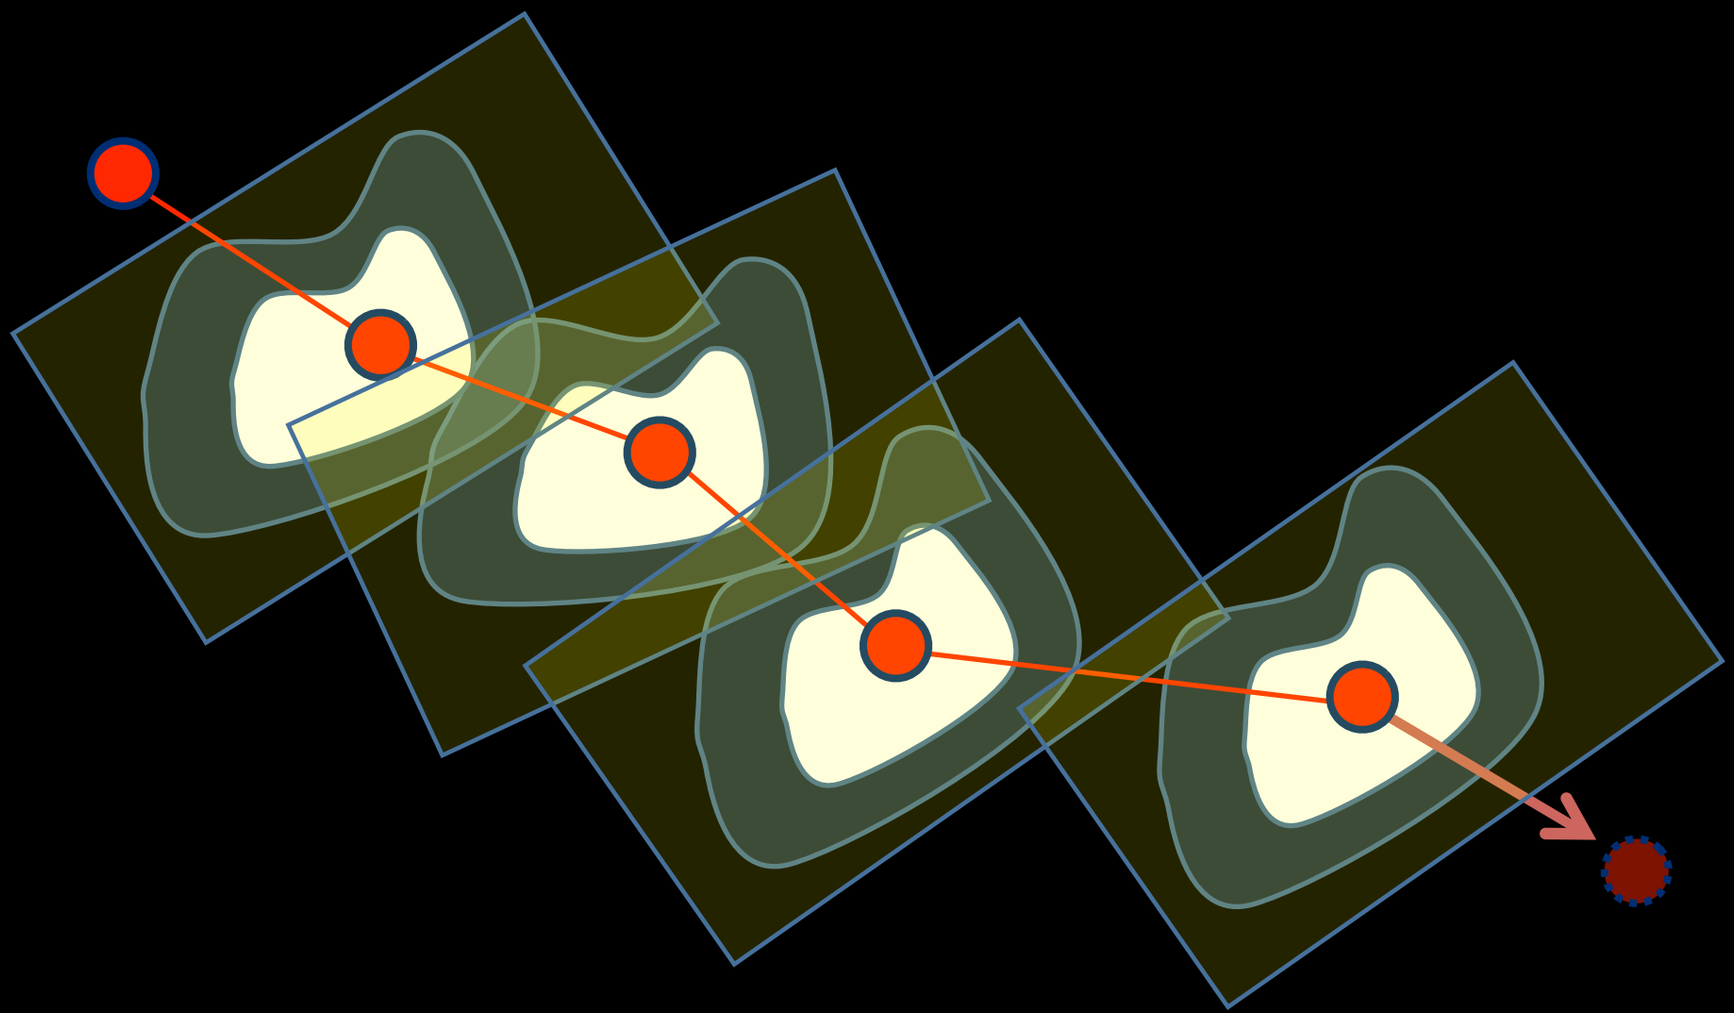

NeuroTrace [Jeong 09]

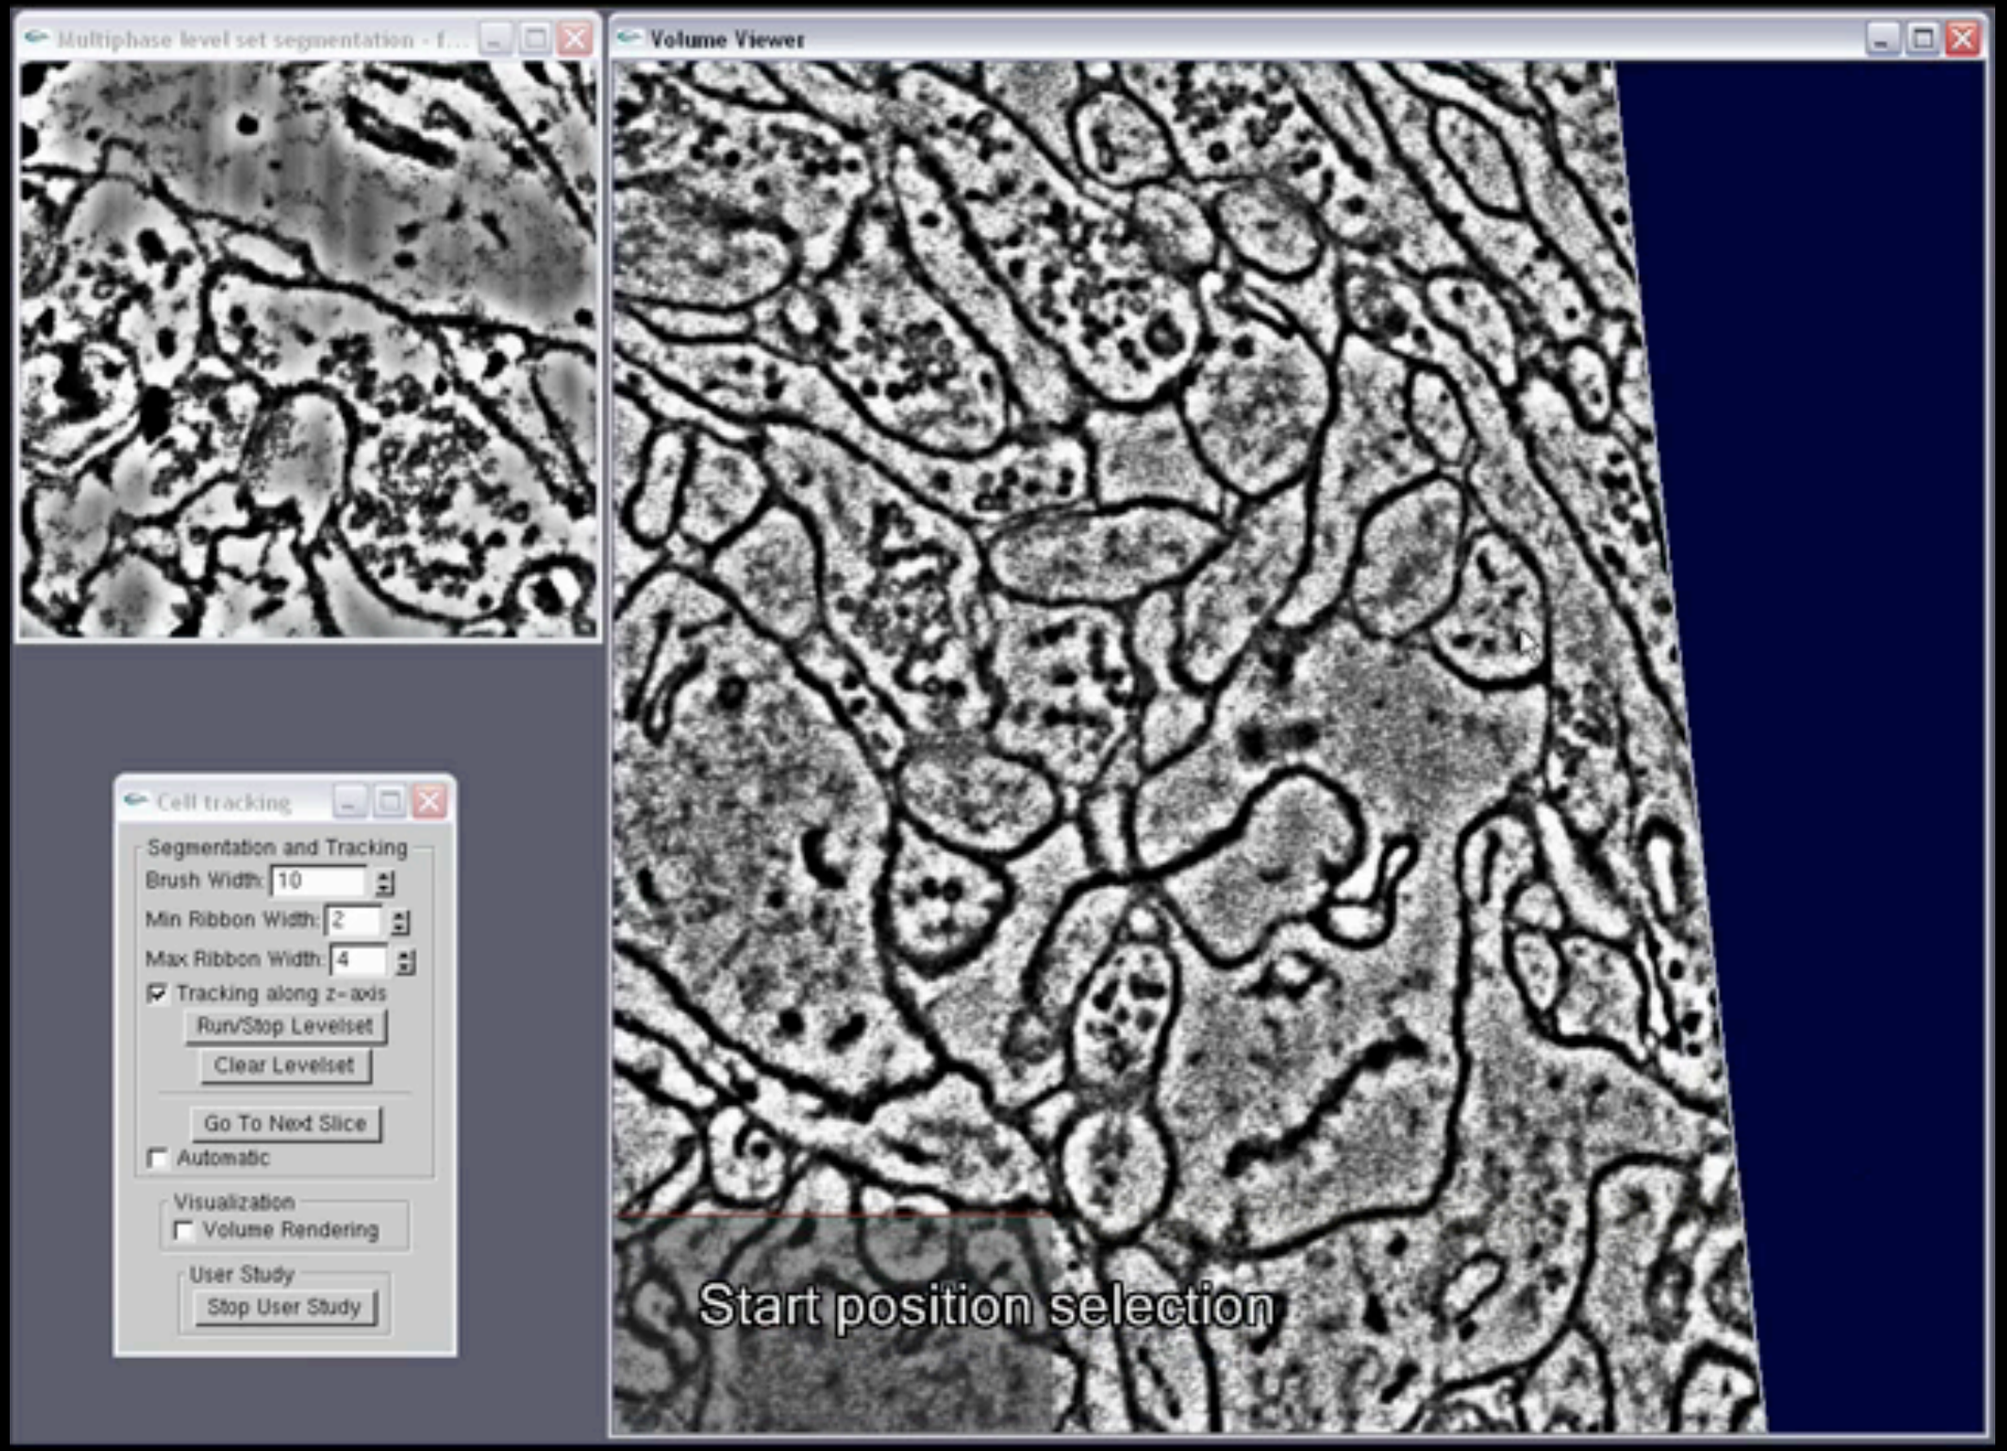

Visualization

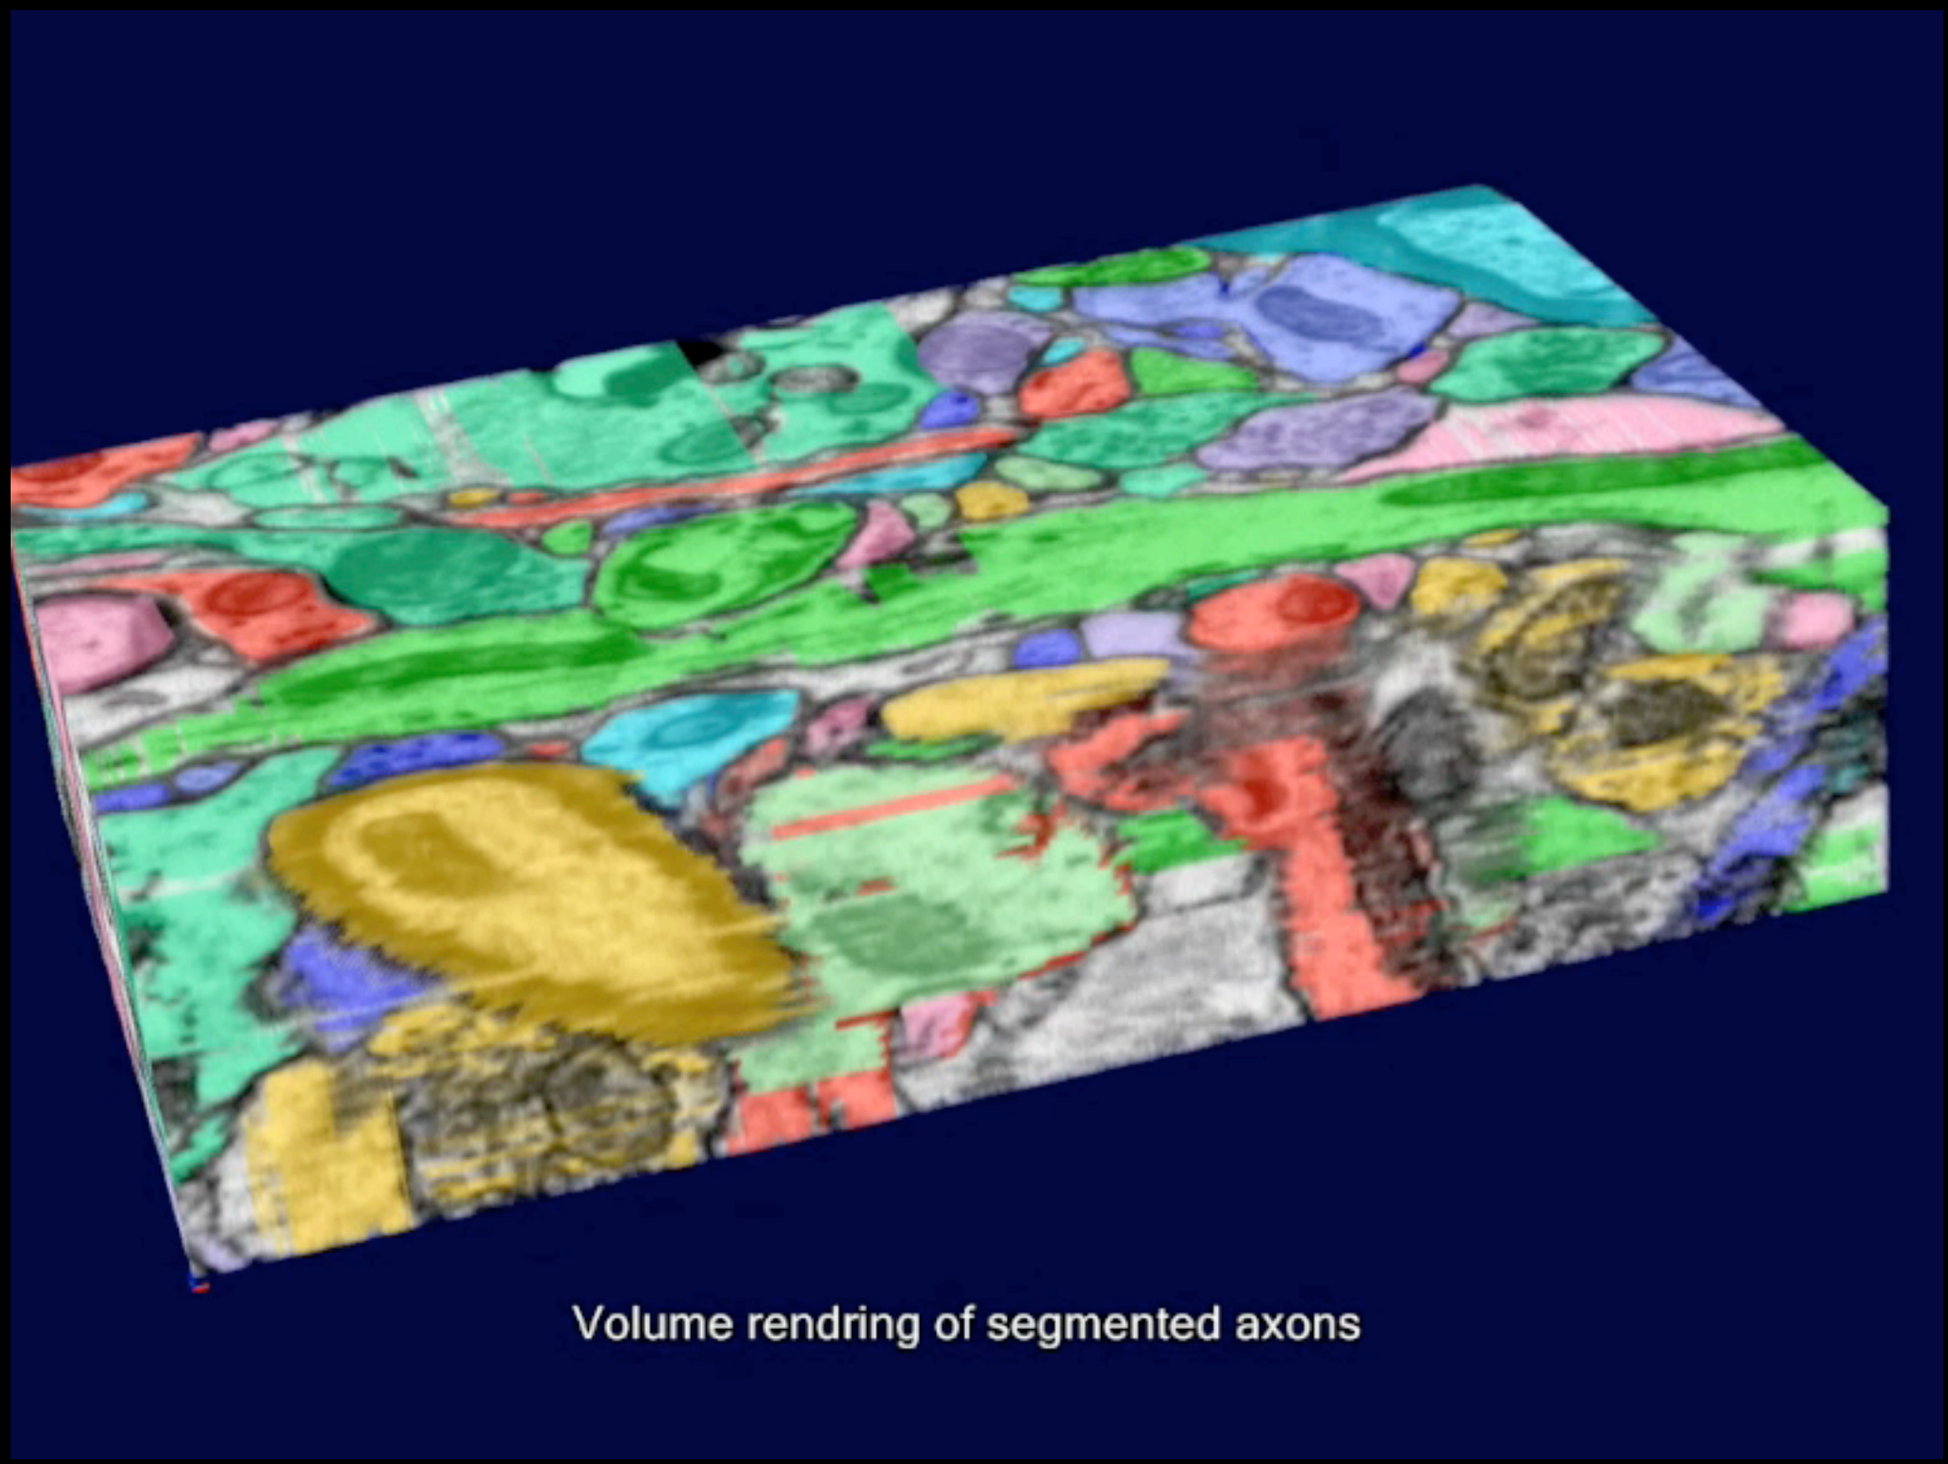

Volume rendring of segmented axons

Visualization

- Elliptical approximation of ribbons

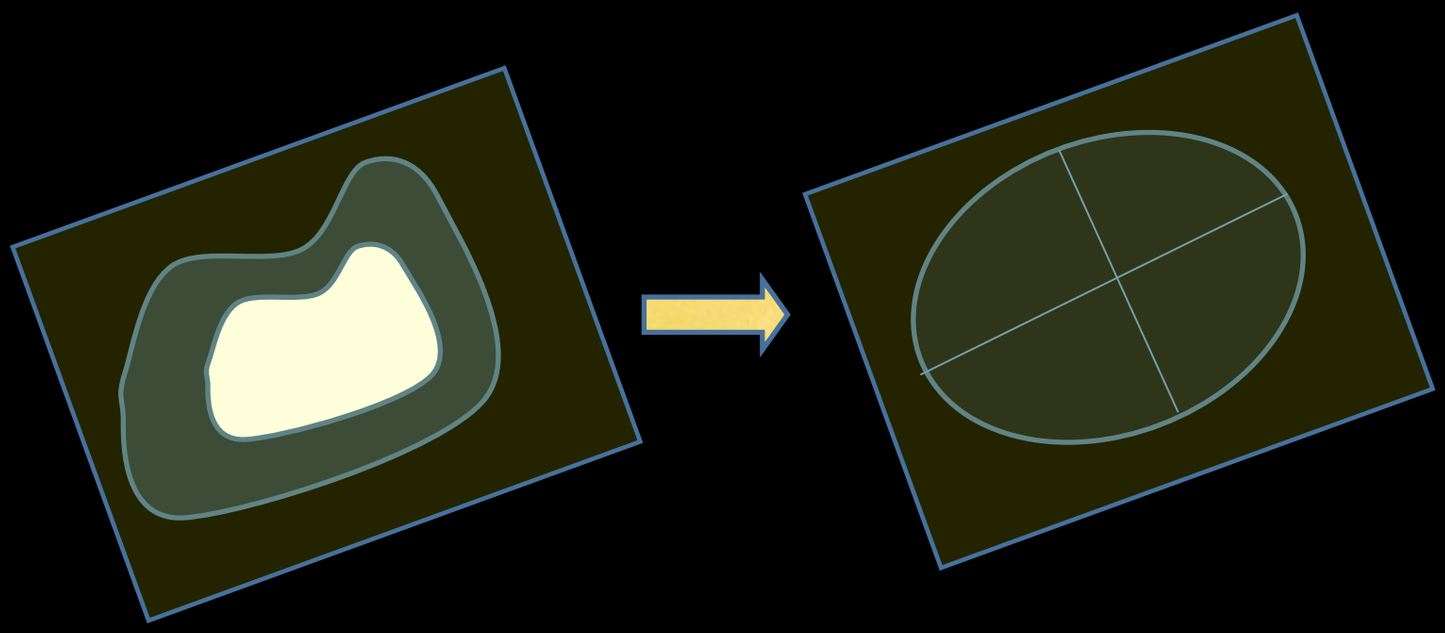

Compression ratio = ~ 1000

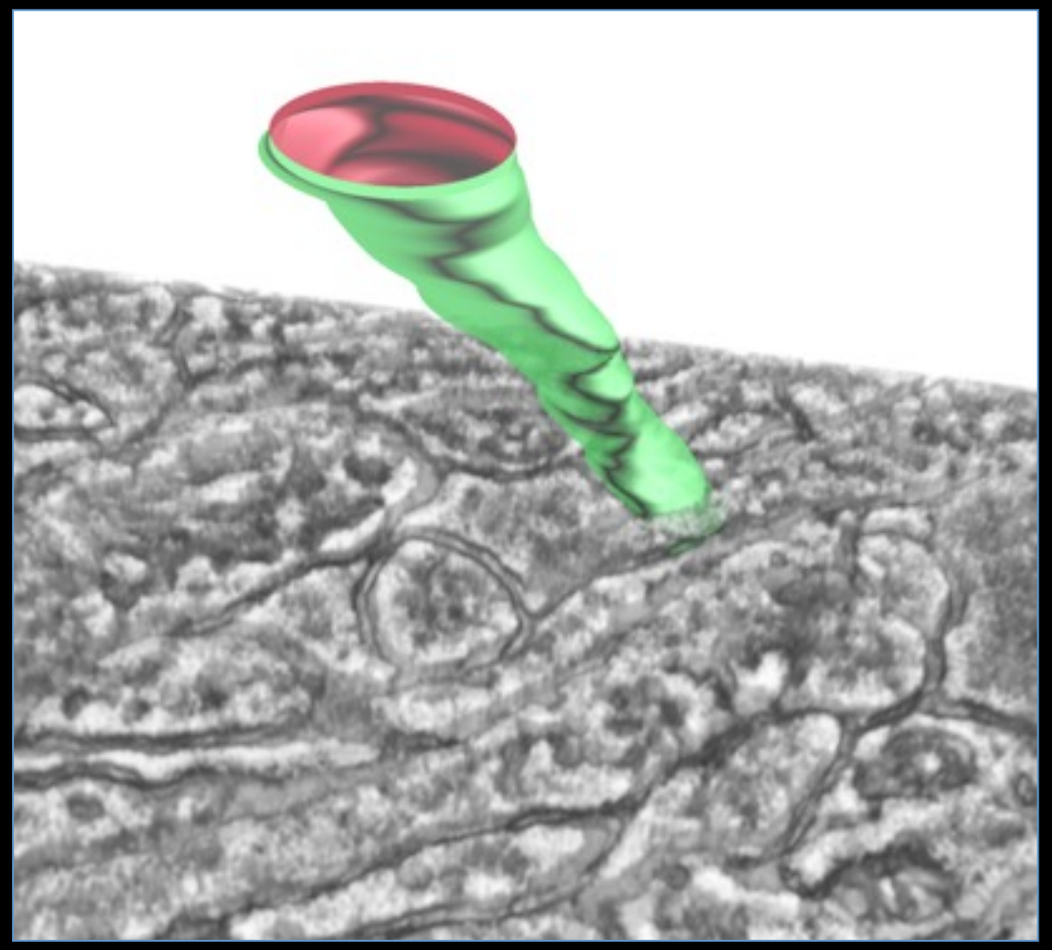

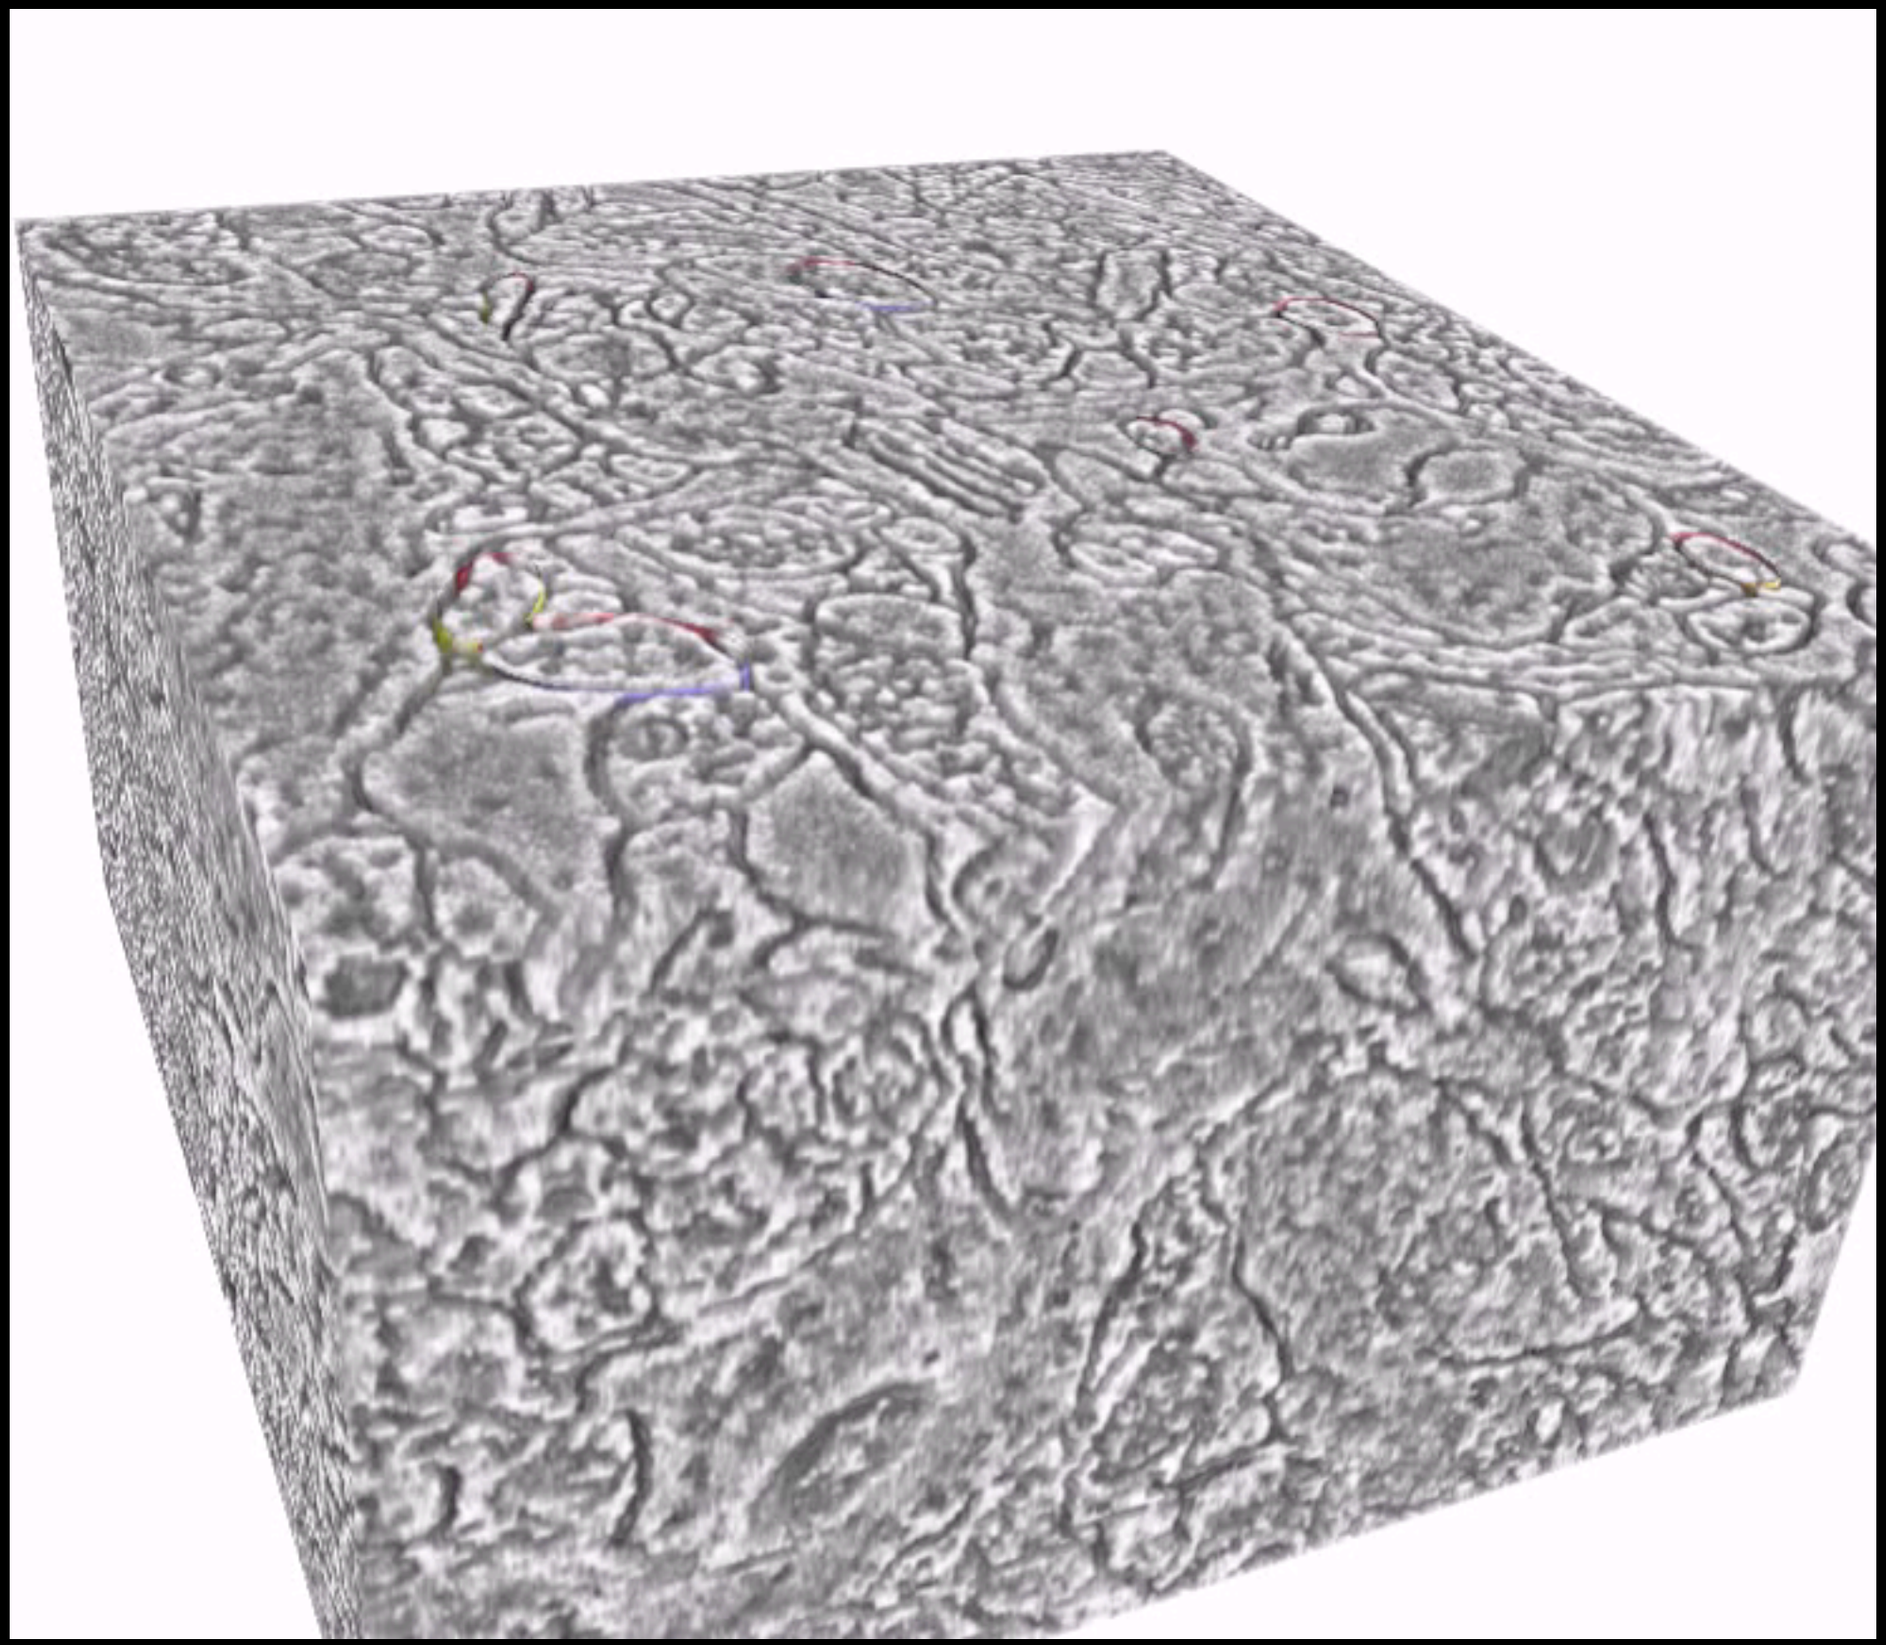

High-Throughput Neuroscience

Select Region

Images

Geometry

EM

Reconstruction

Visualization
Analysis

Multi-Scale Imaging

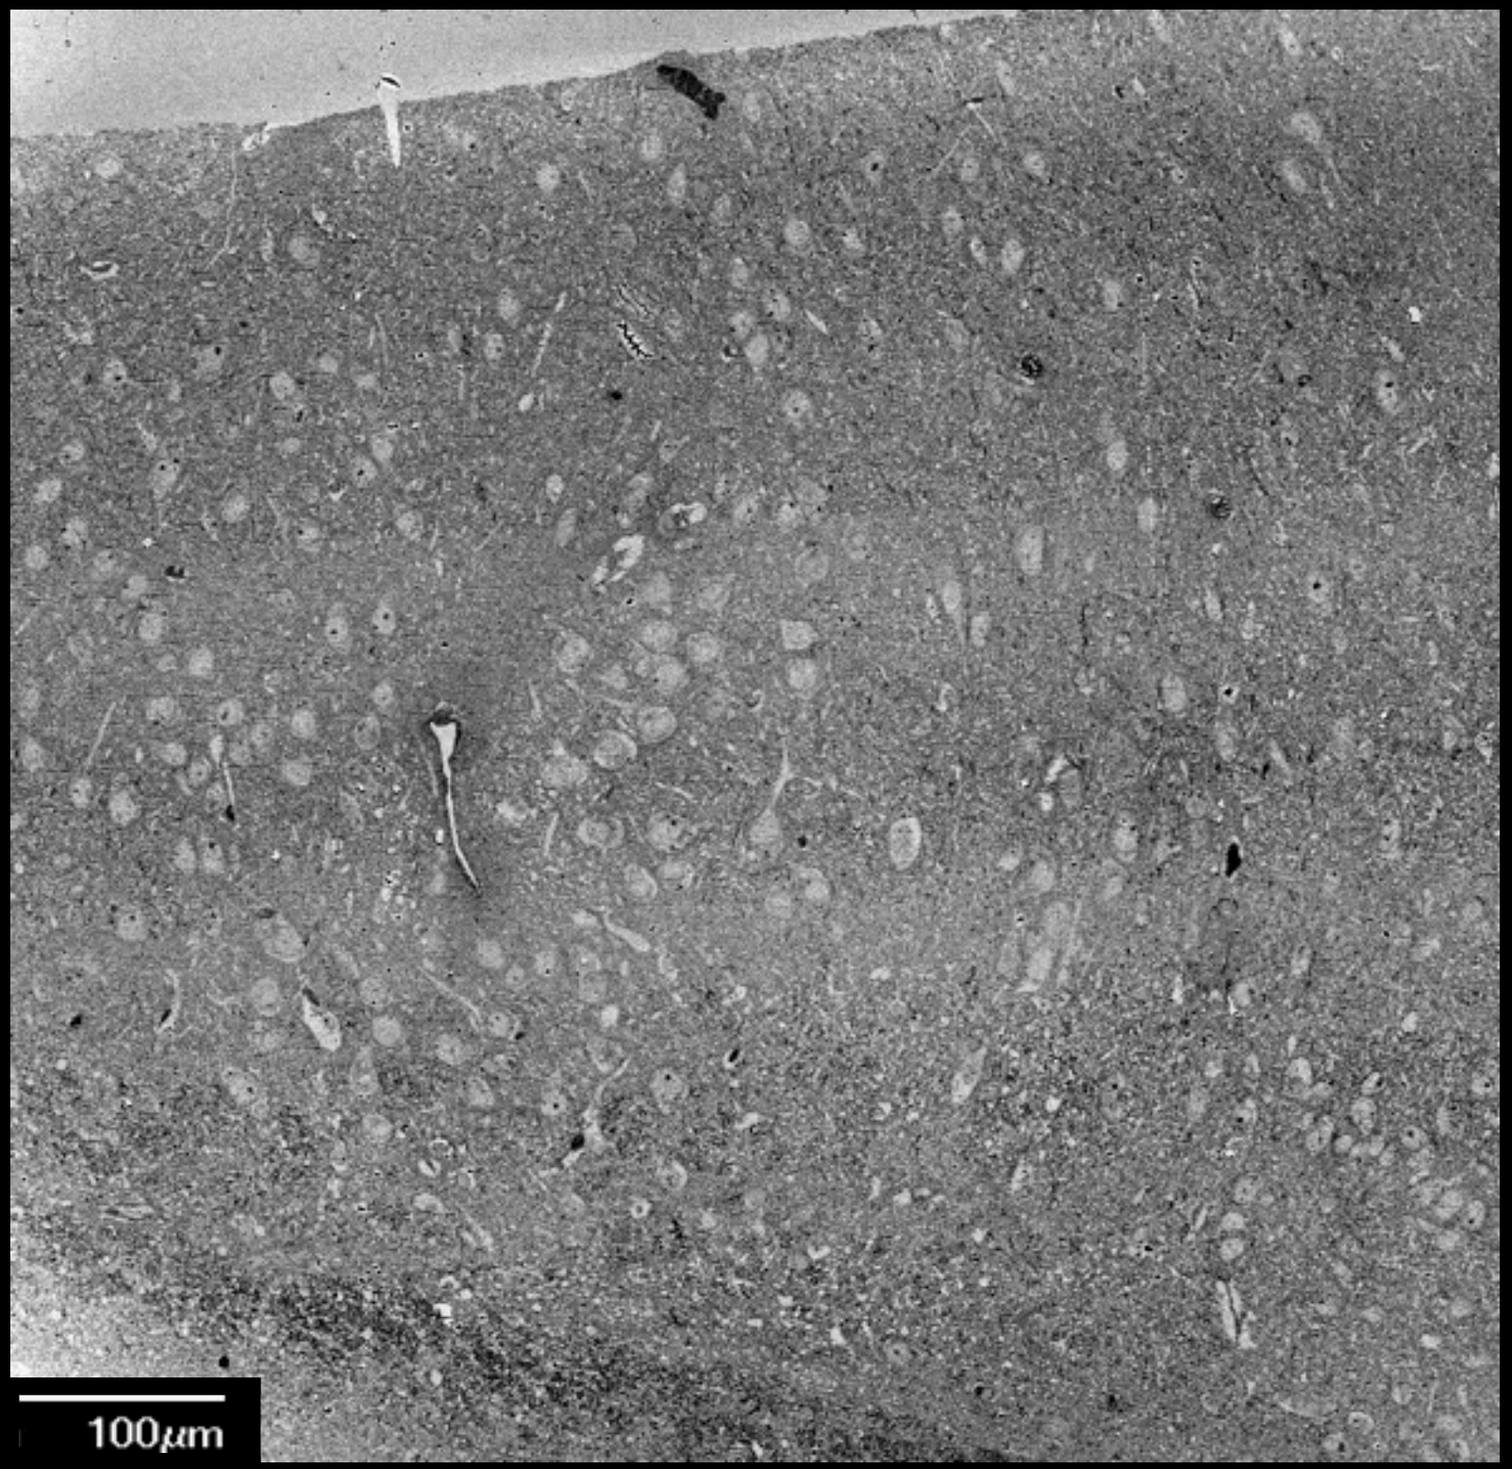

Multi-Scale Imaging

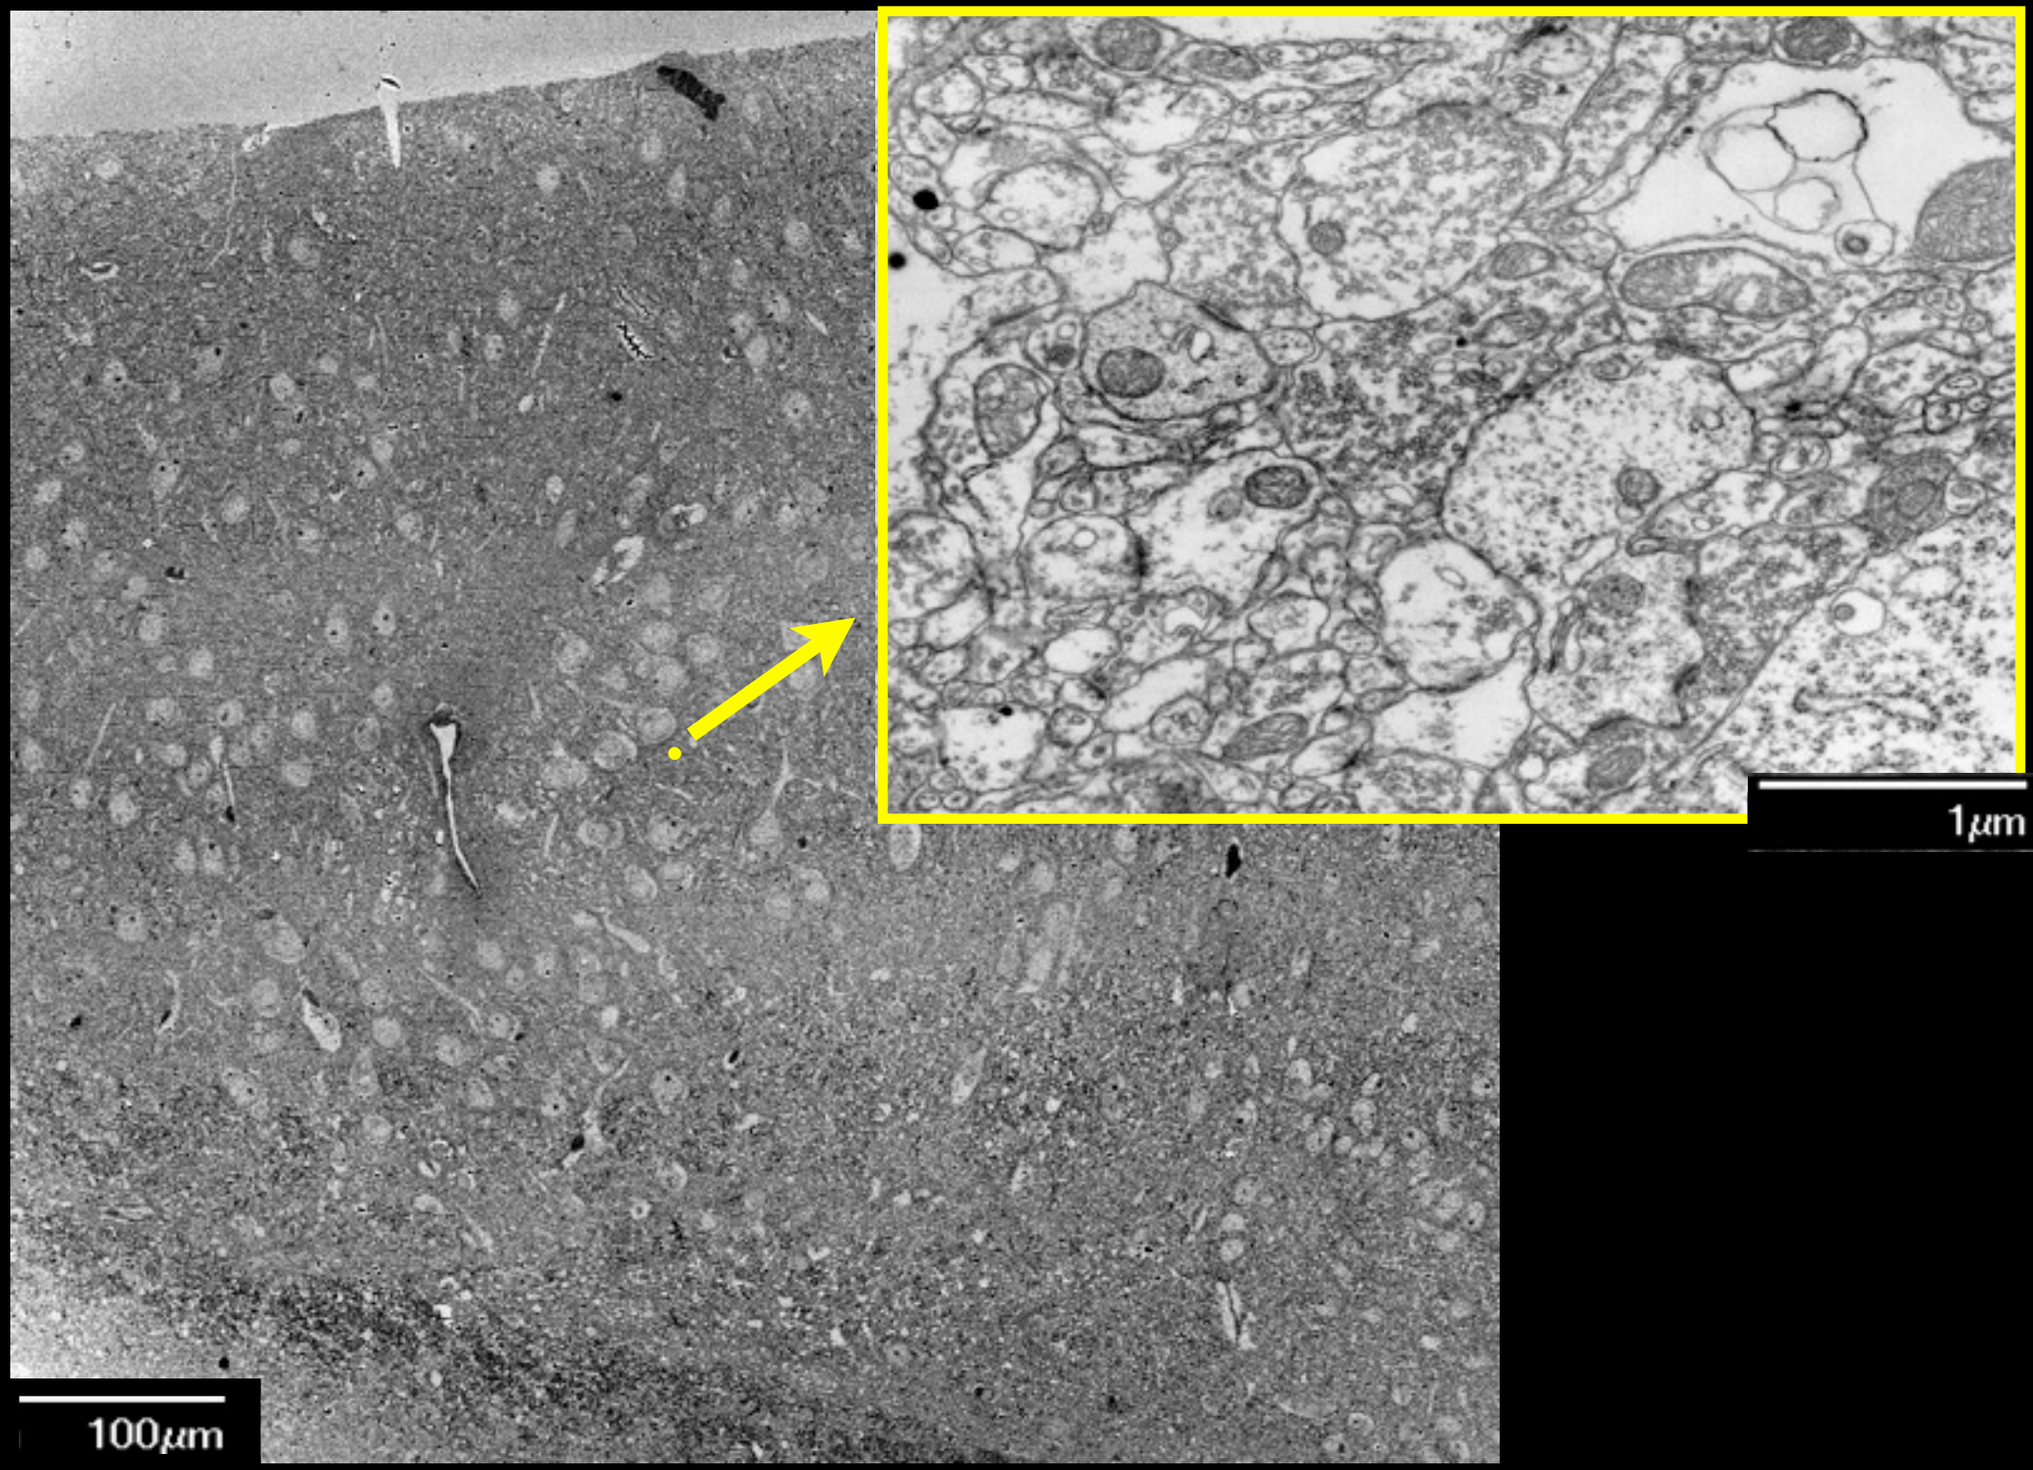

Multi-Scale Imaging

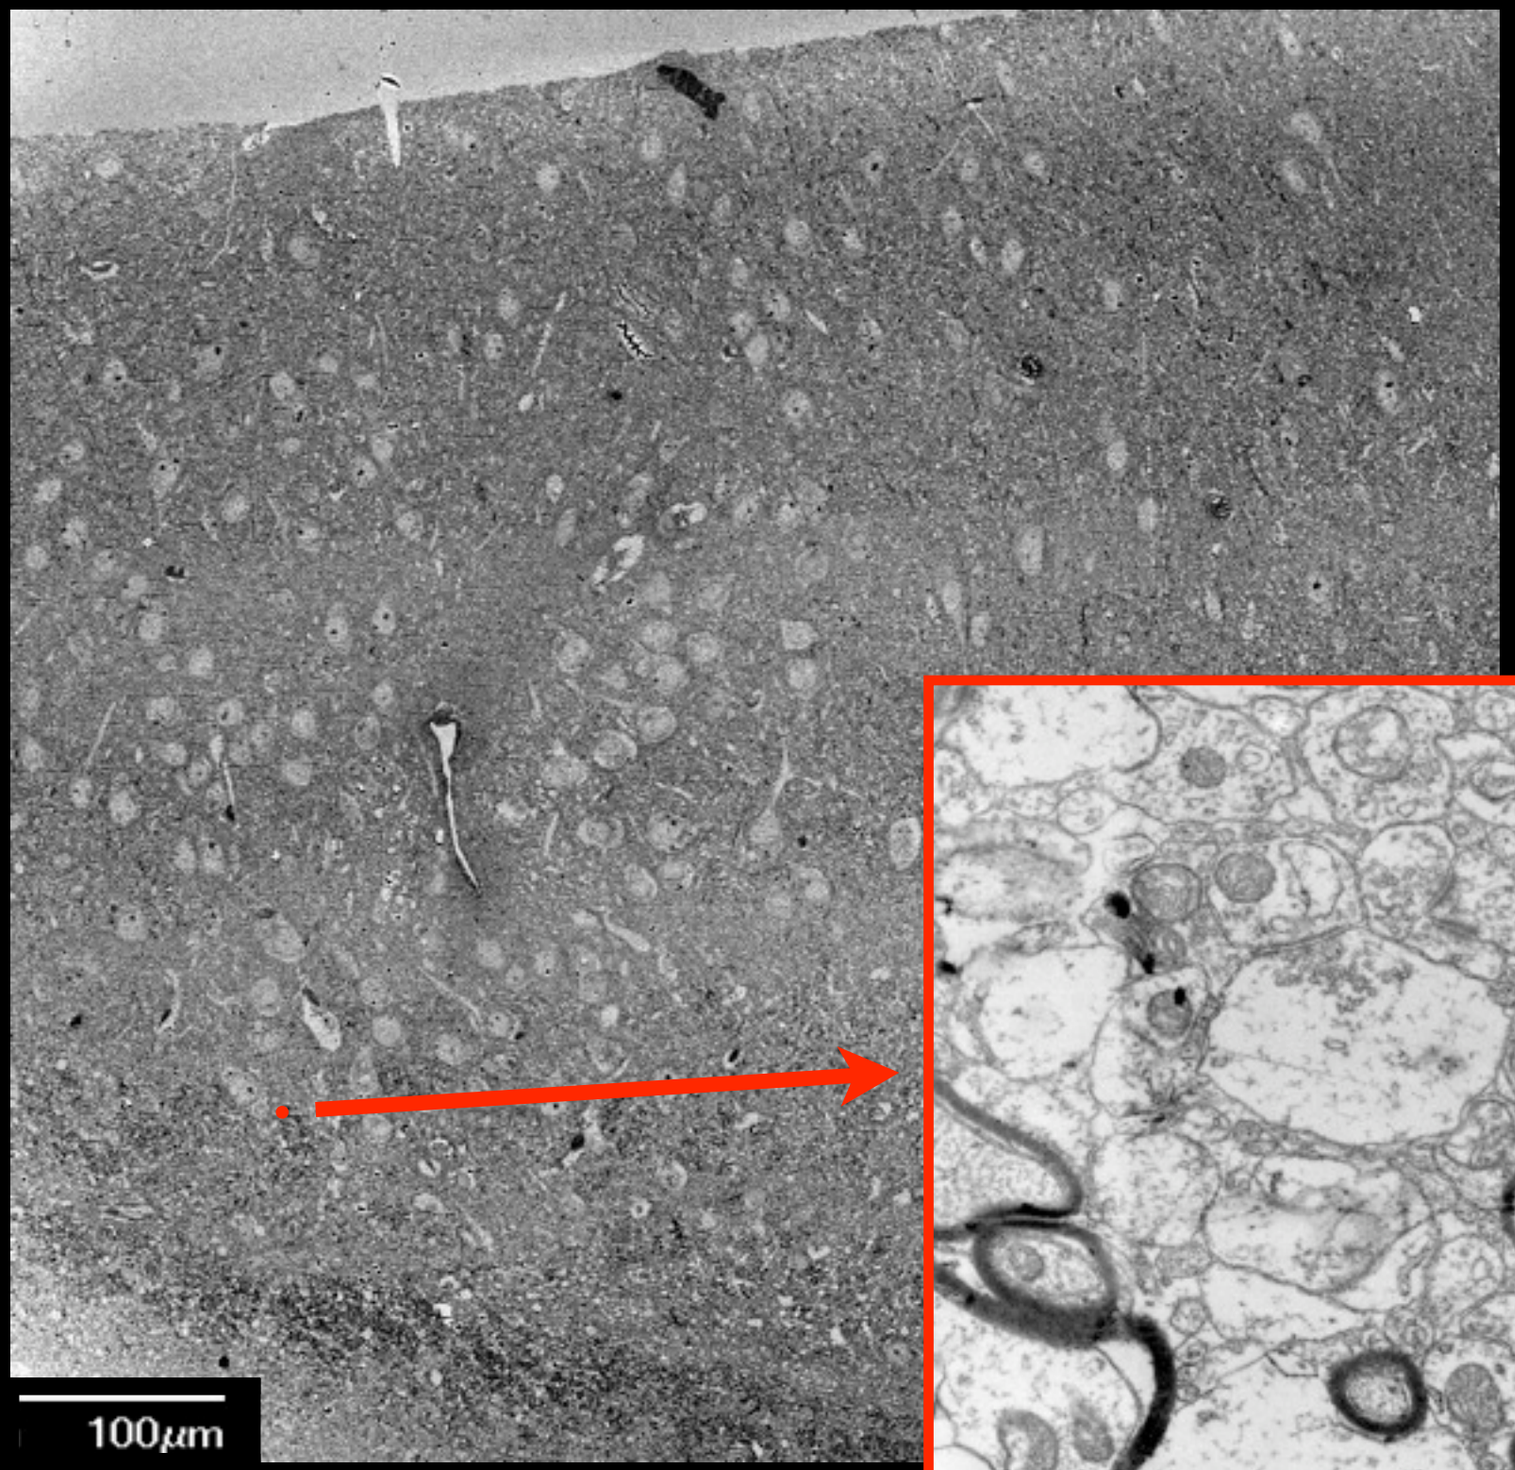

Scalable

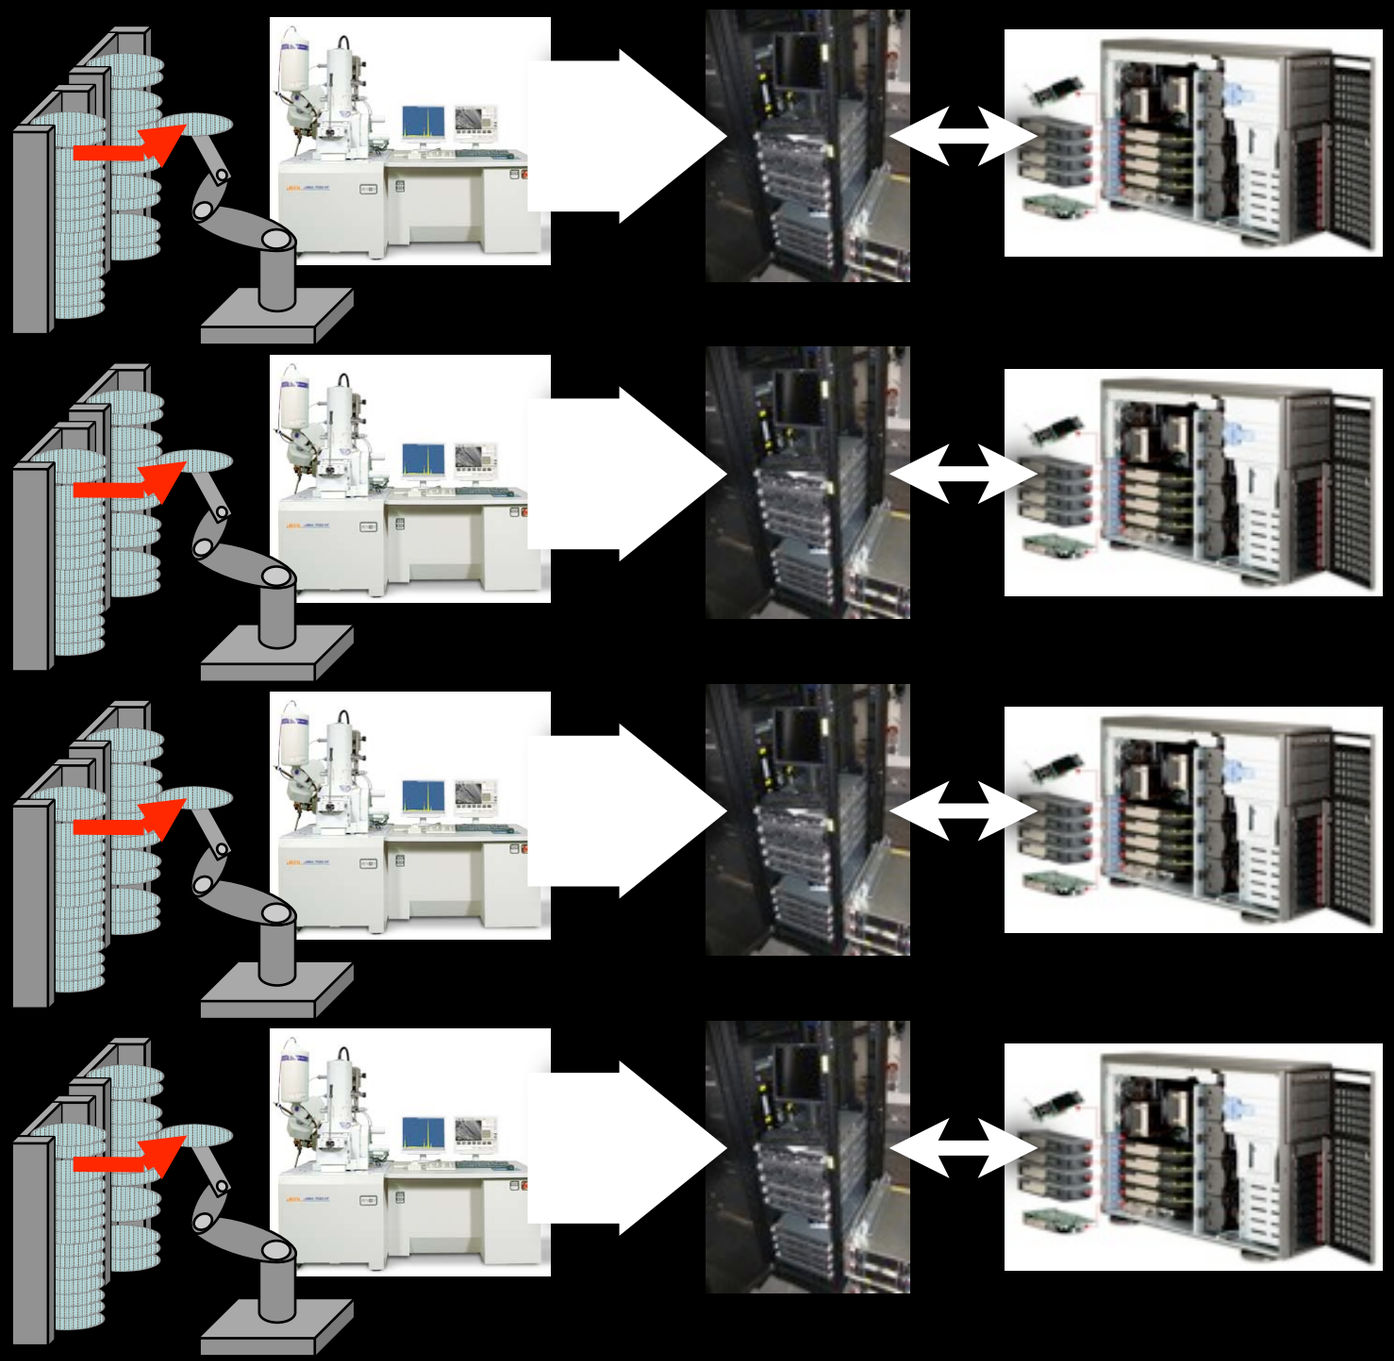

The Future

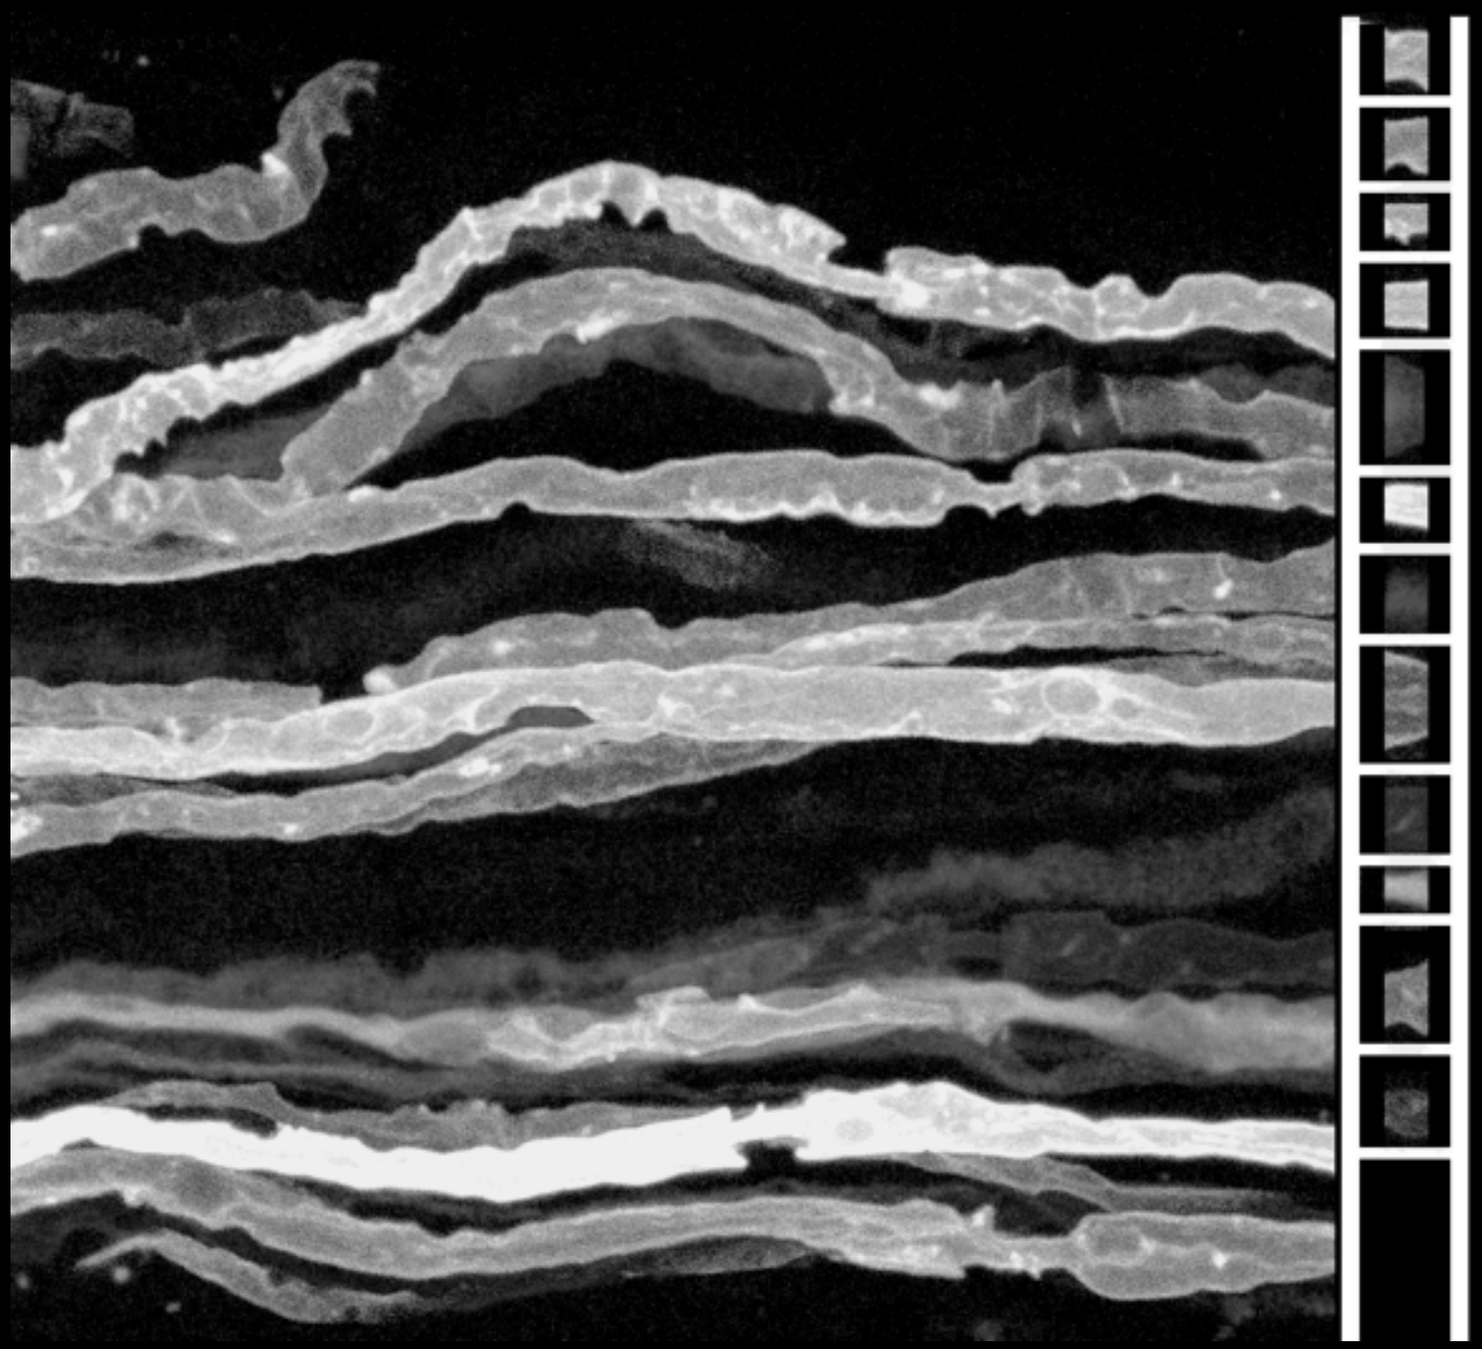

J. Livet

Brainbow

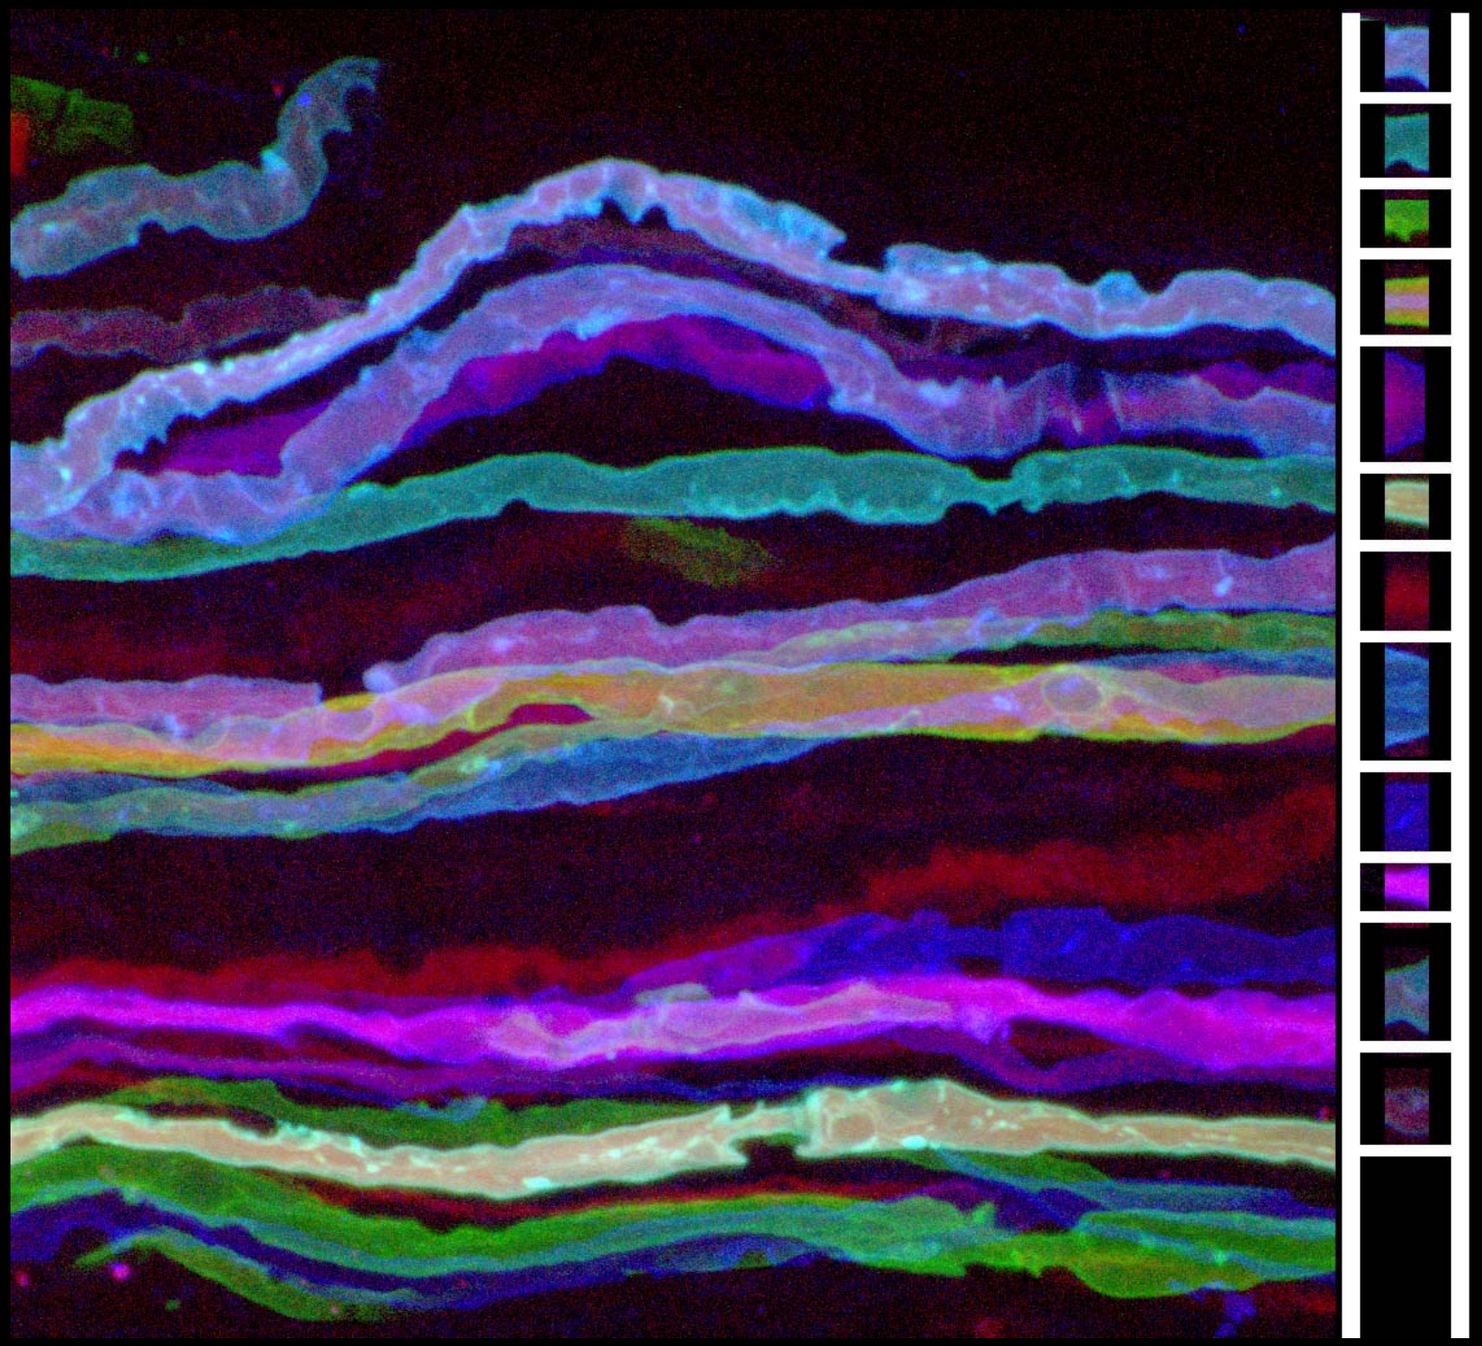

J. Livet

Cerebellum

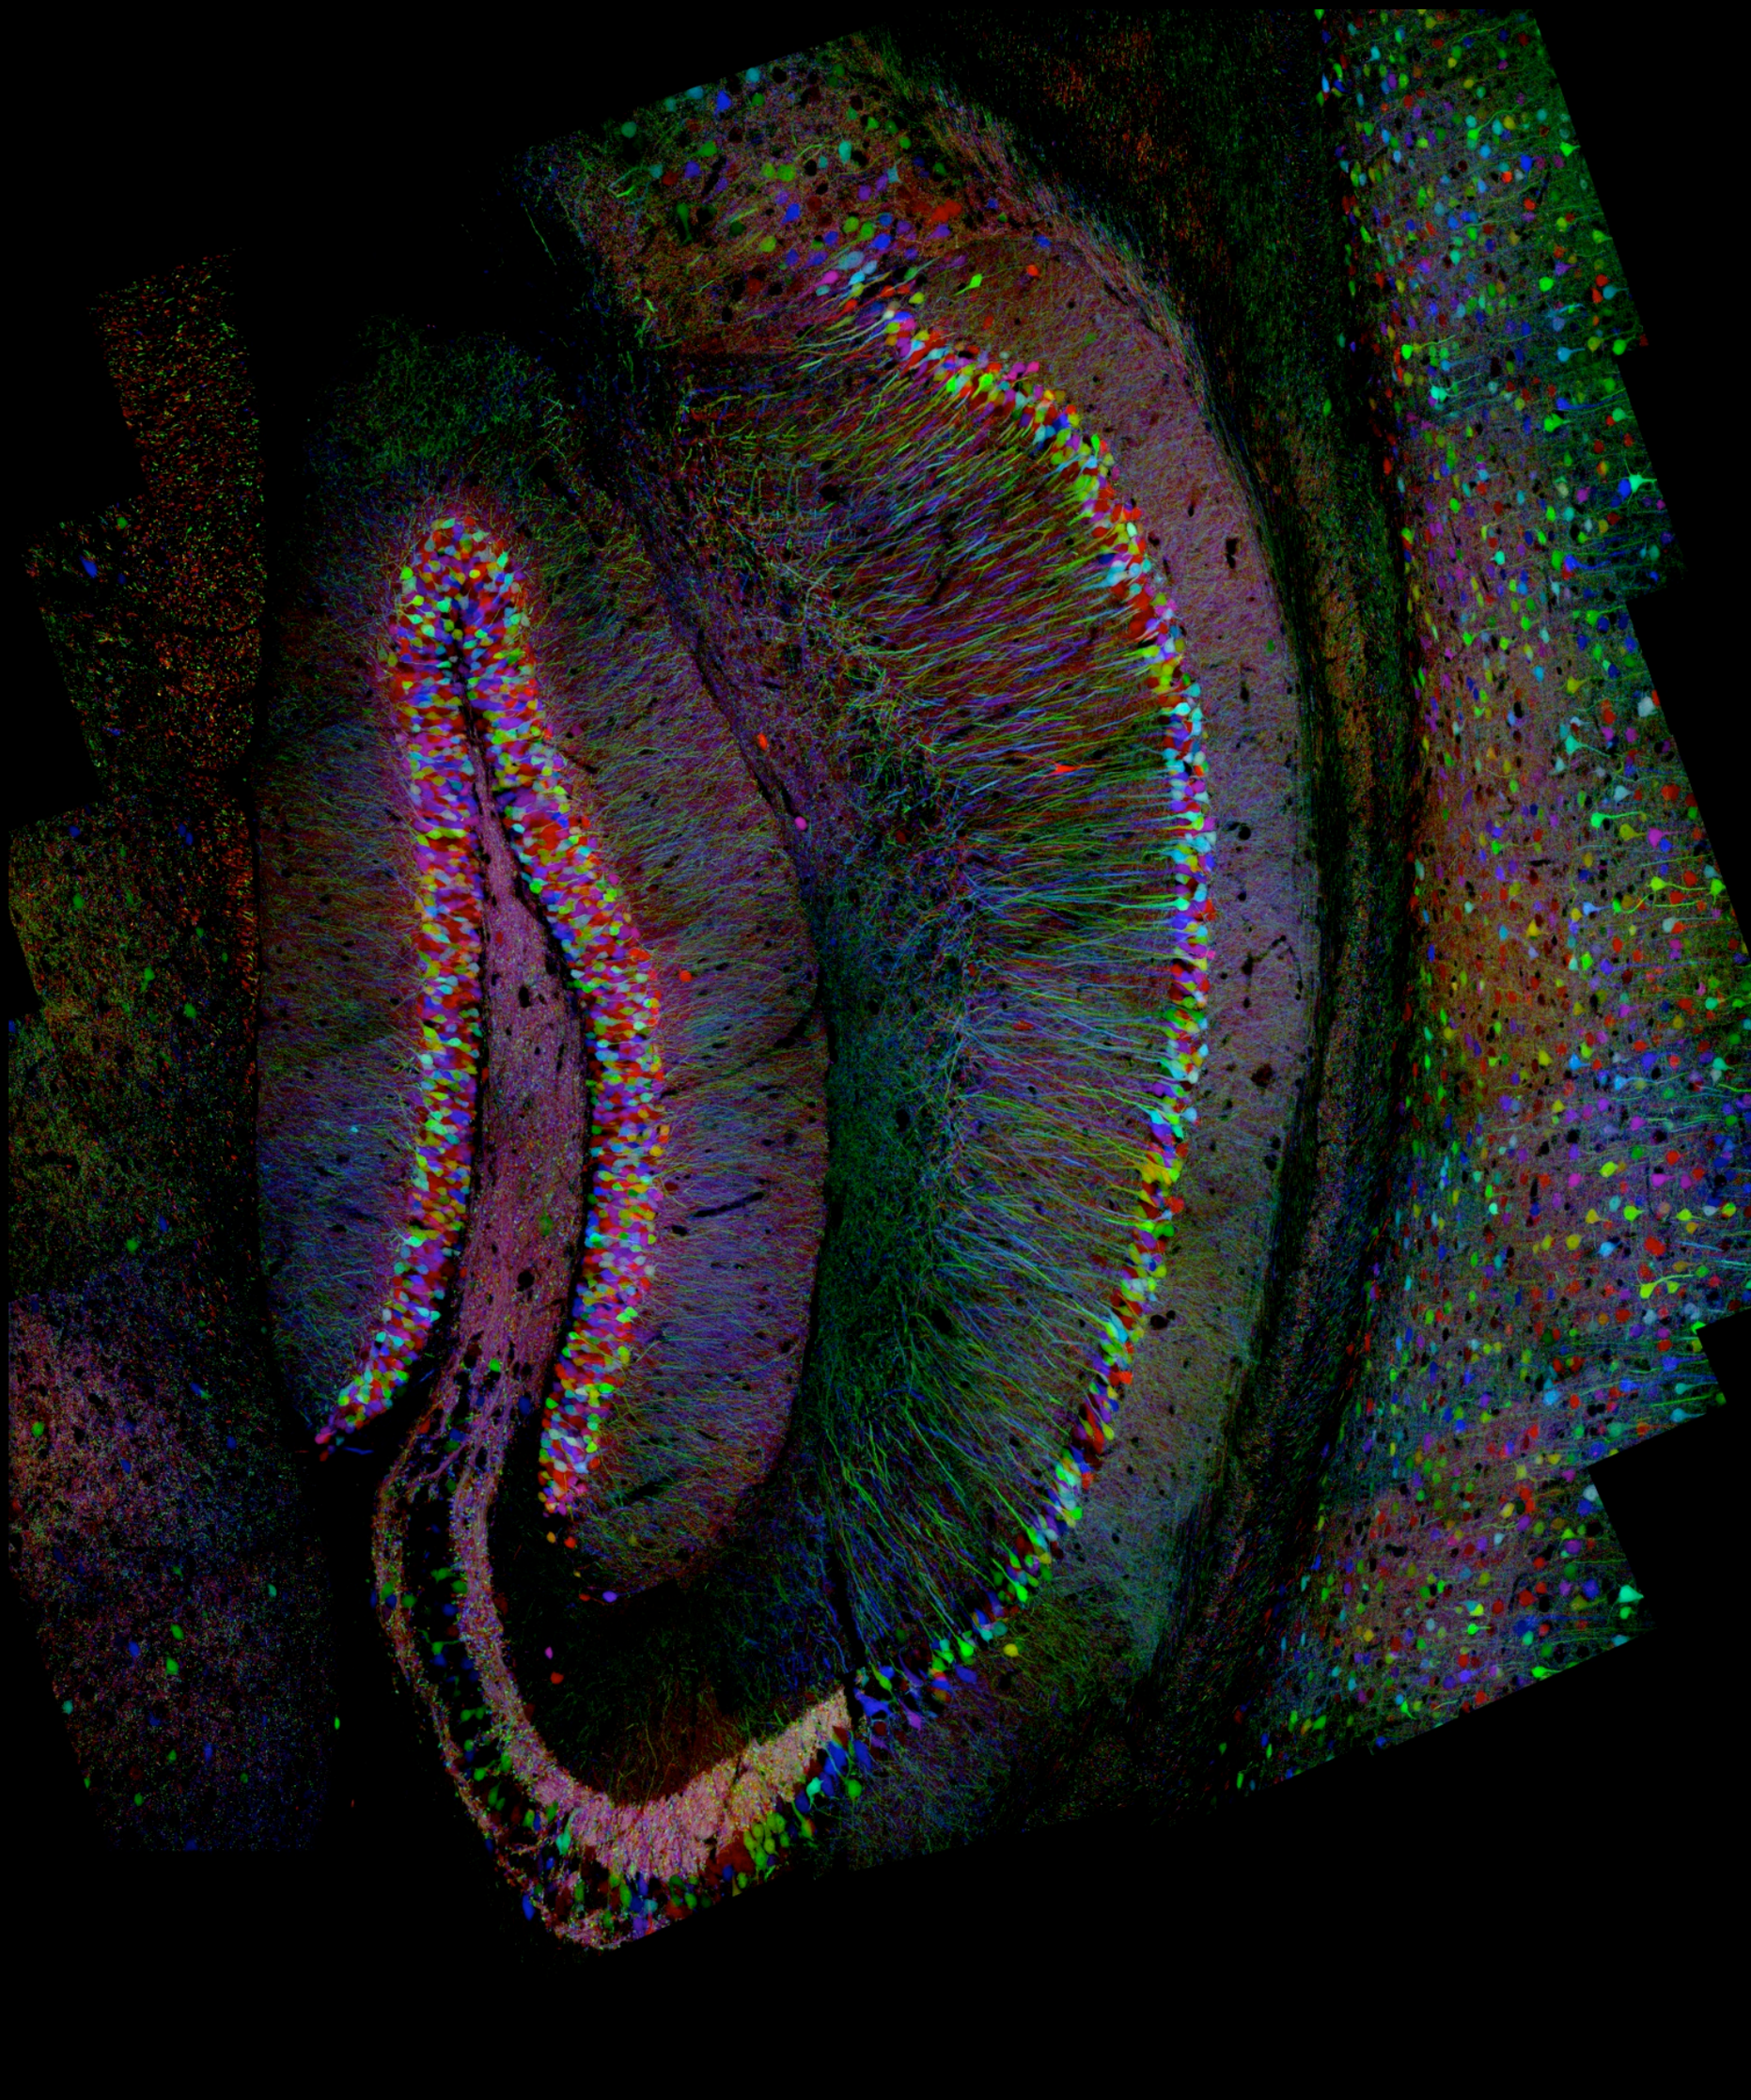

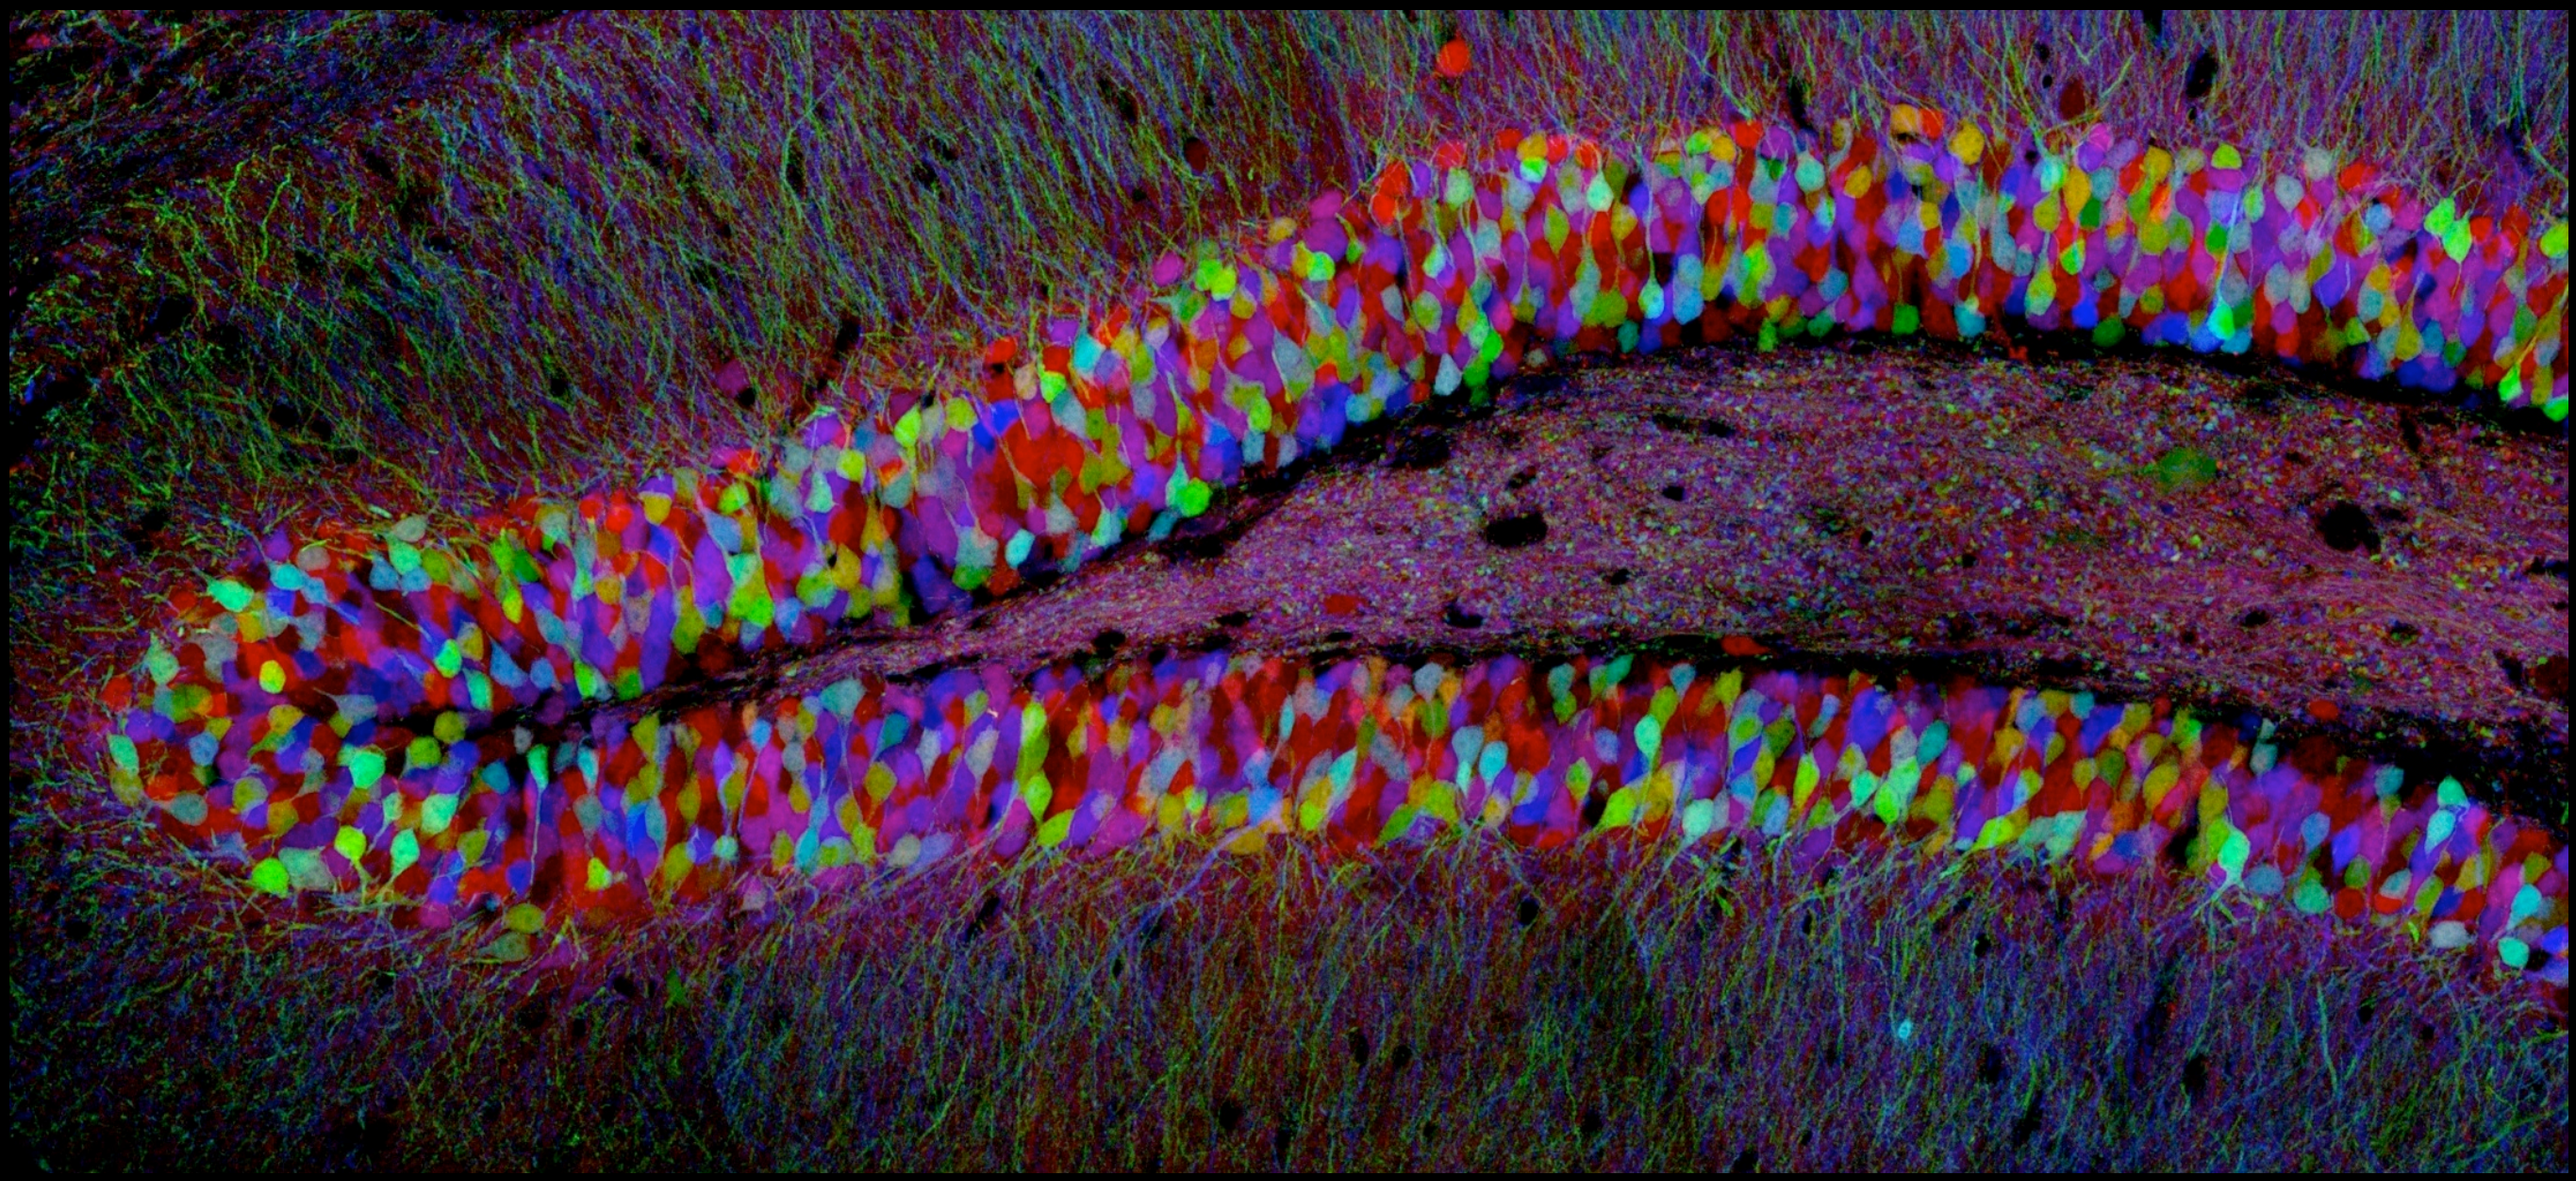

Hippocampus

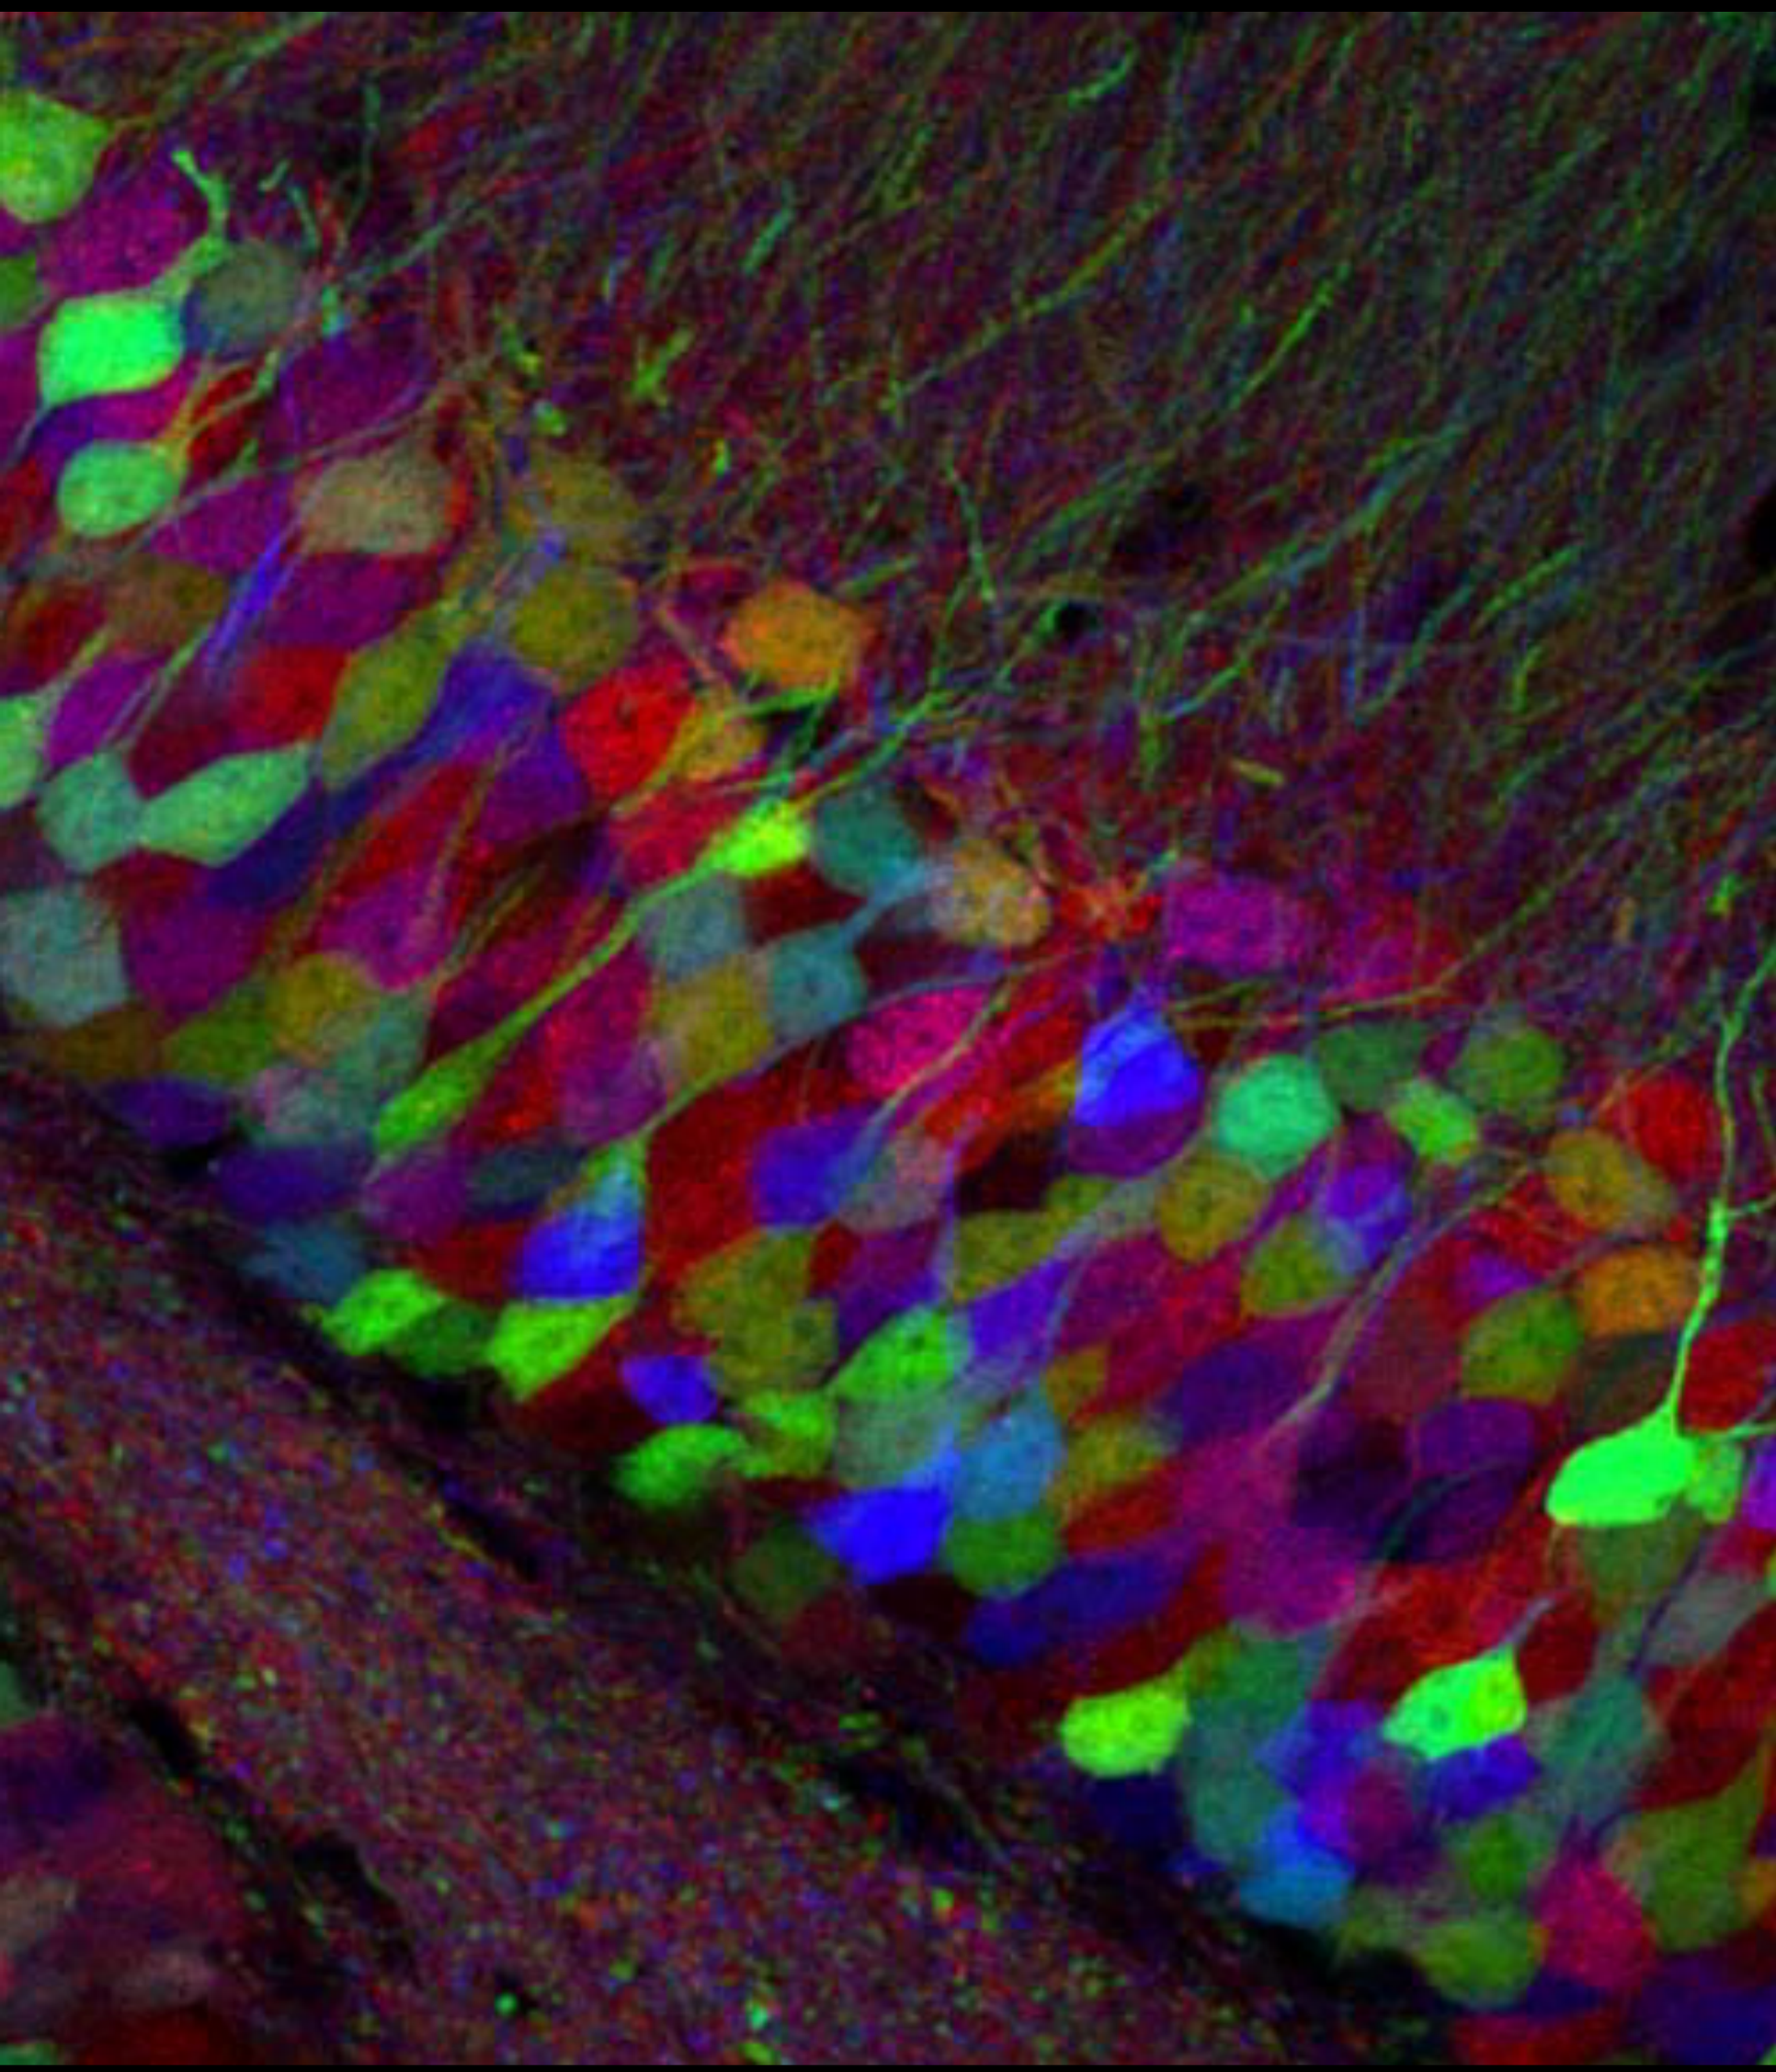

Convergent evolution?

More Information

Won-Ki Jeong
Reconstructing the Brain: Extracting Neural Circuitry with CUDA and MPI
Session Id#1075
Friday 4 pm Empire

The MWA

Outback Supercomputing

Collaborators

- Harvard Center for Astrophysics
  - Lincoln Greenhill
  - Daniel Mitchell
  - Stephen Ord
  - Randall Wayth
- SEAS
  - Kevin Dale, Richard Edgar
  - Hanspeter Pfister

The Scientific Challenge

- What happened through the 1 billion years of the universe's dark ages?

~300,000 ~1 bio

3D Tomography

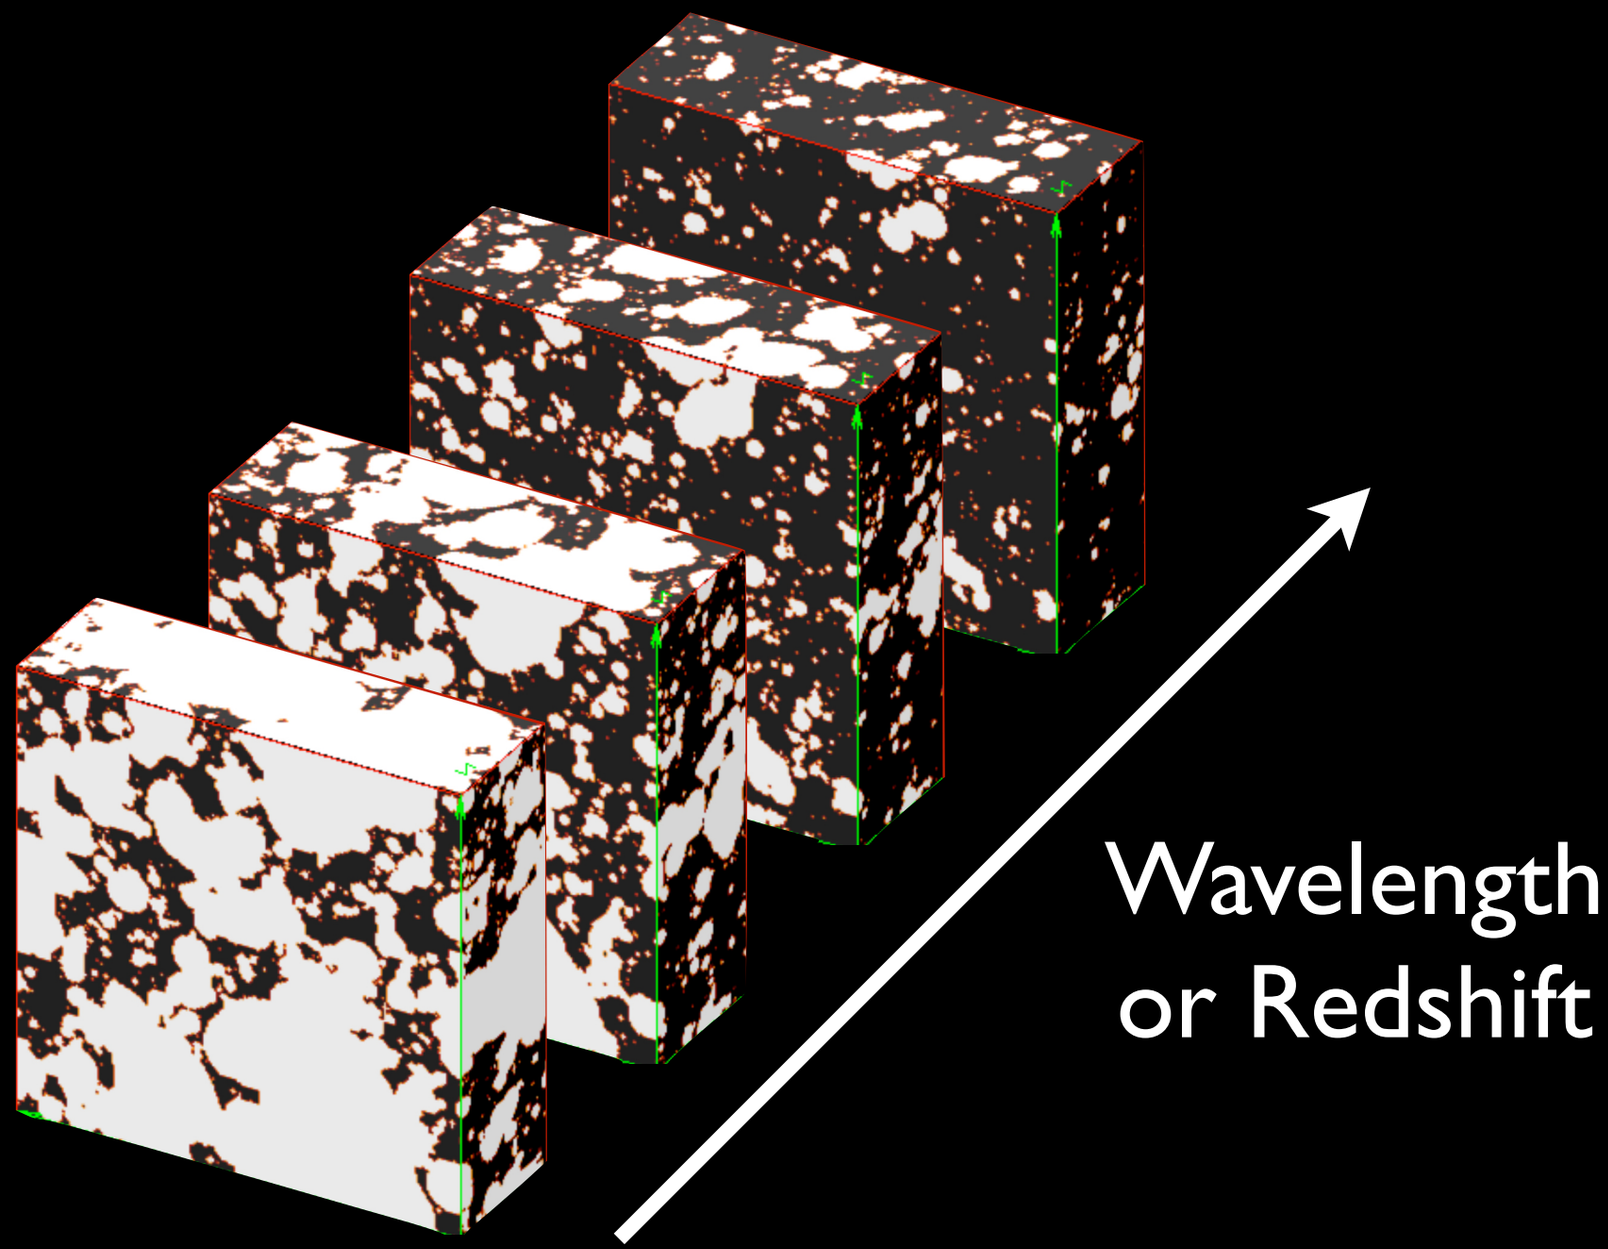

Radio Interferometry

Correlator

Image in Fourier Domain

From this...

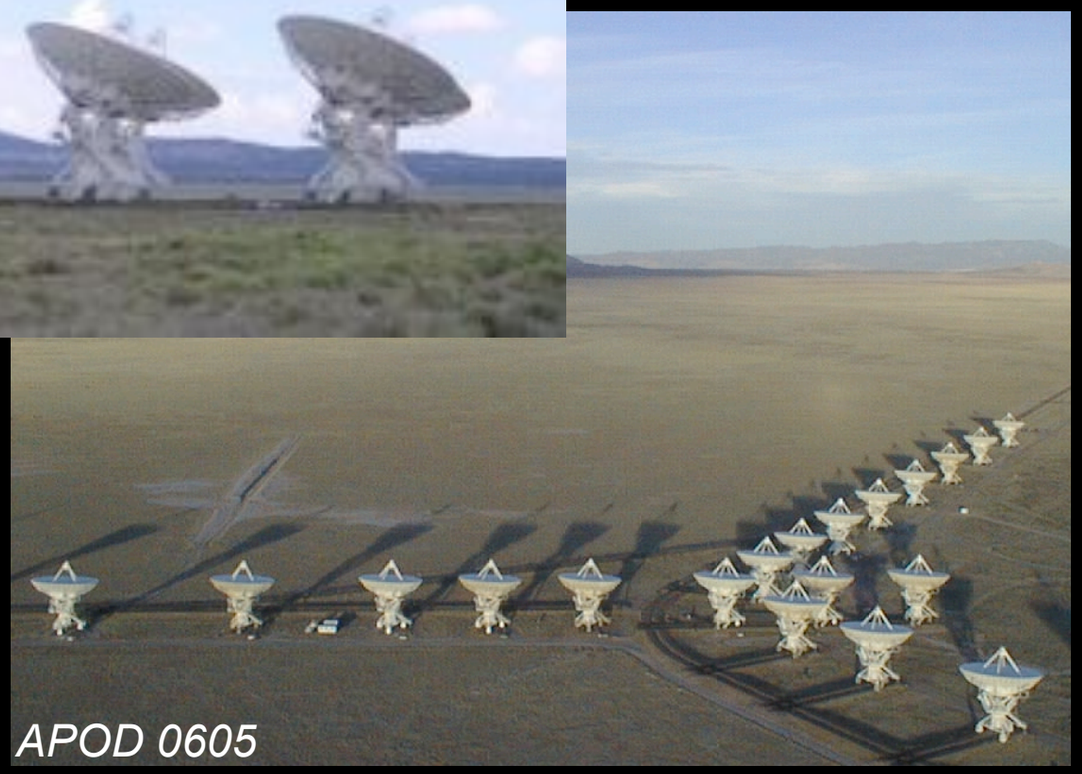

APOD 0605

...to this

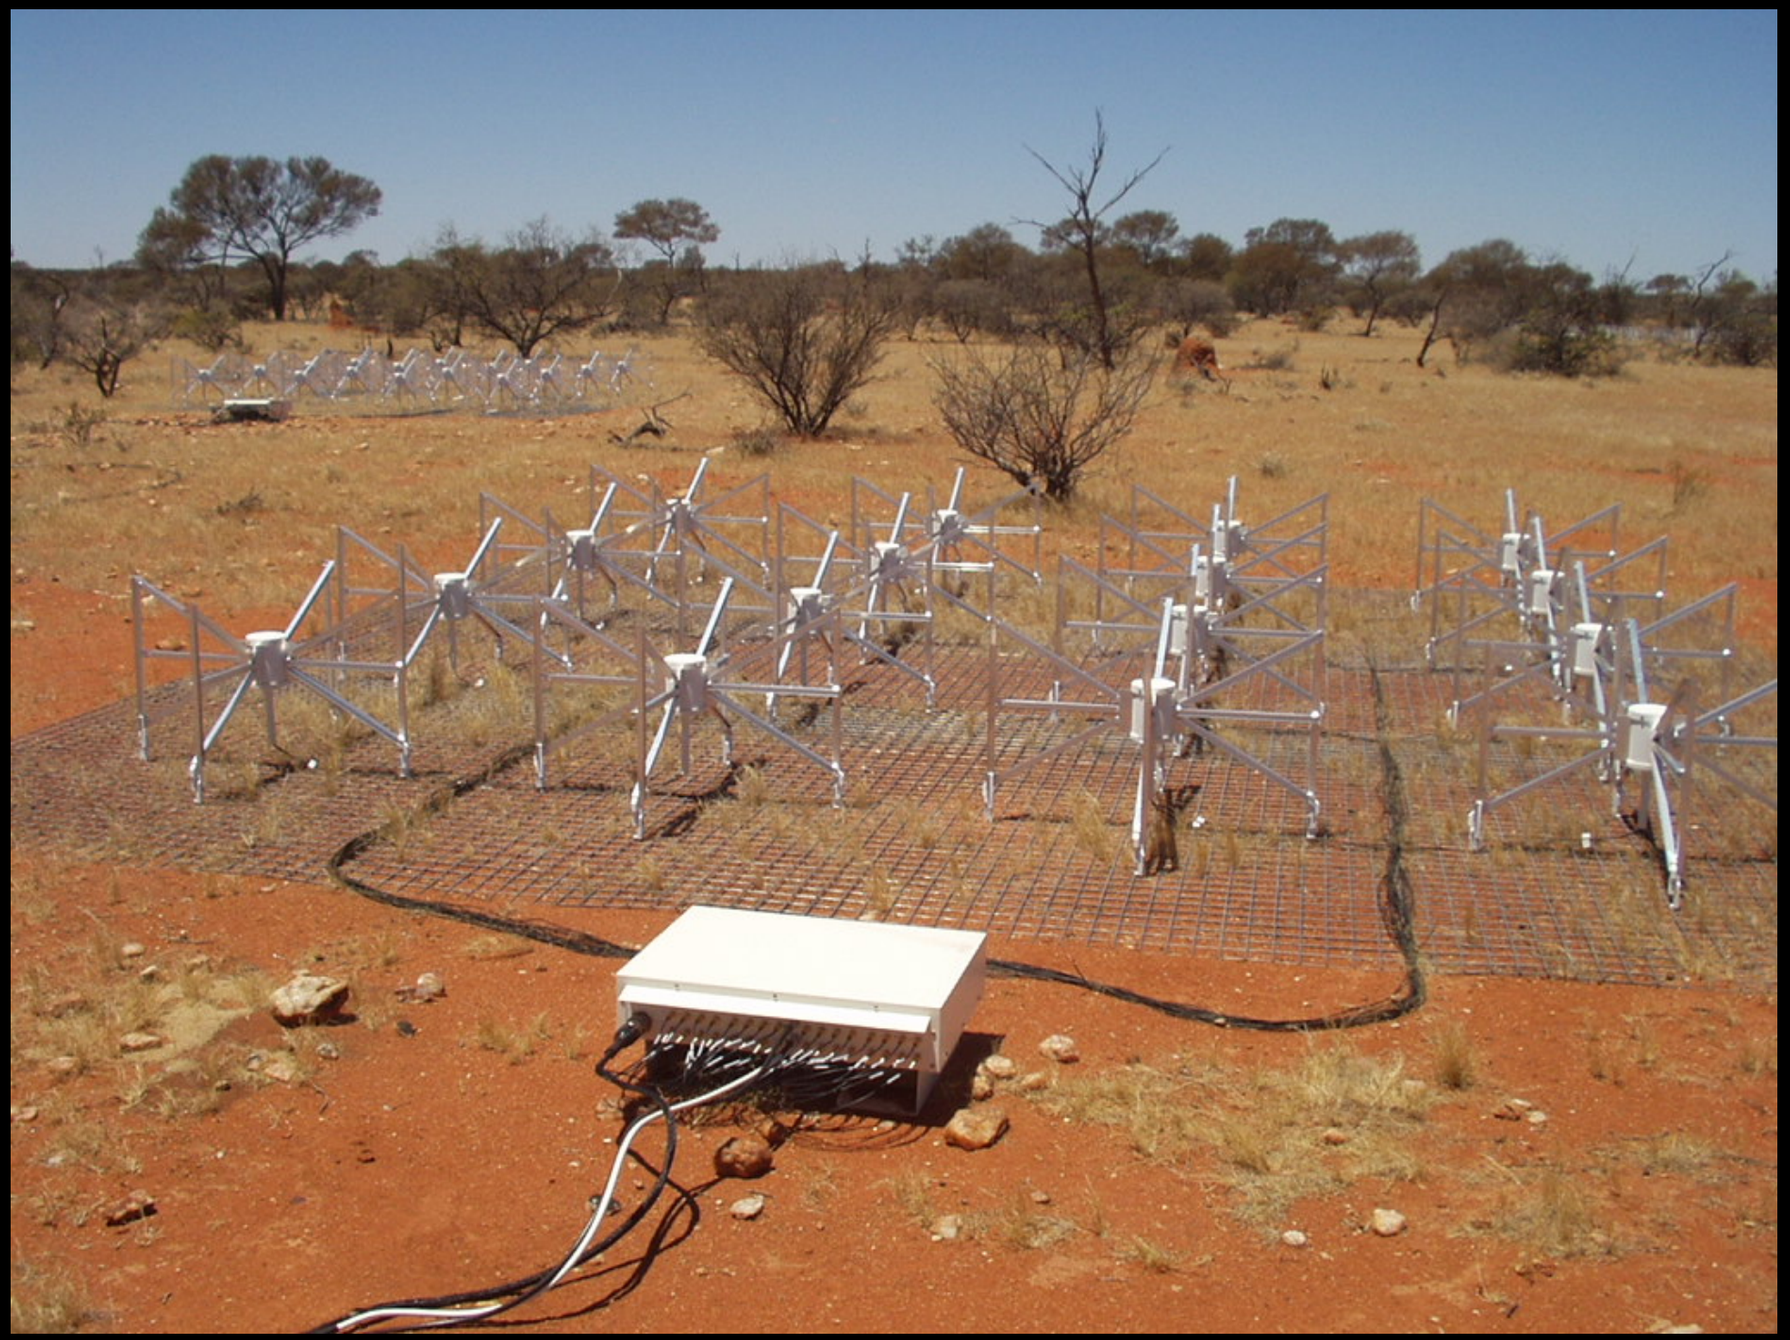

MWA

- 32 antenna prototype
- Eventually 512 antennas on 1 km^2

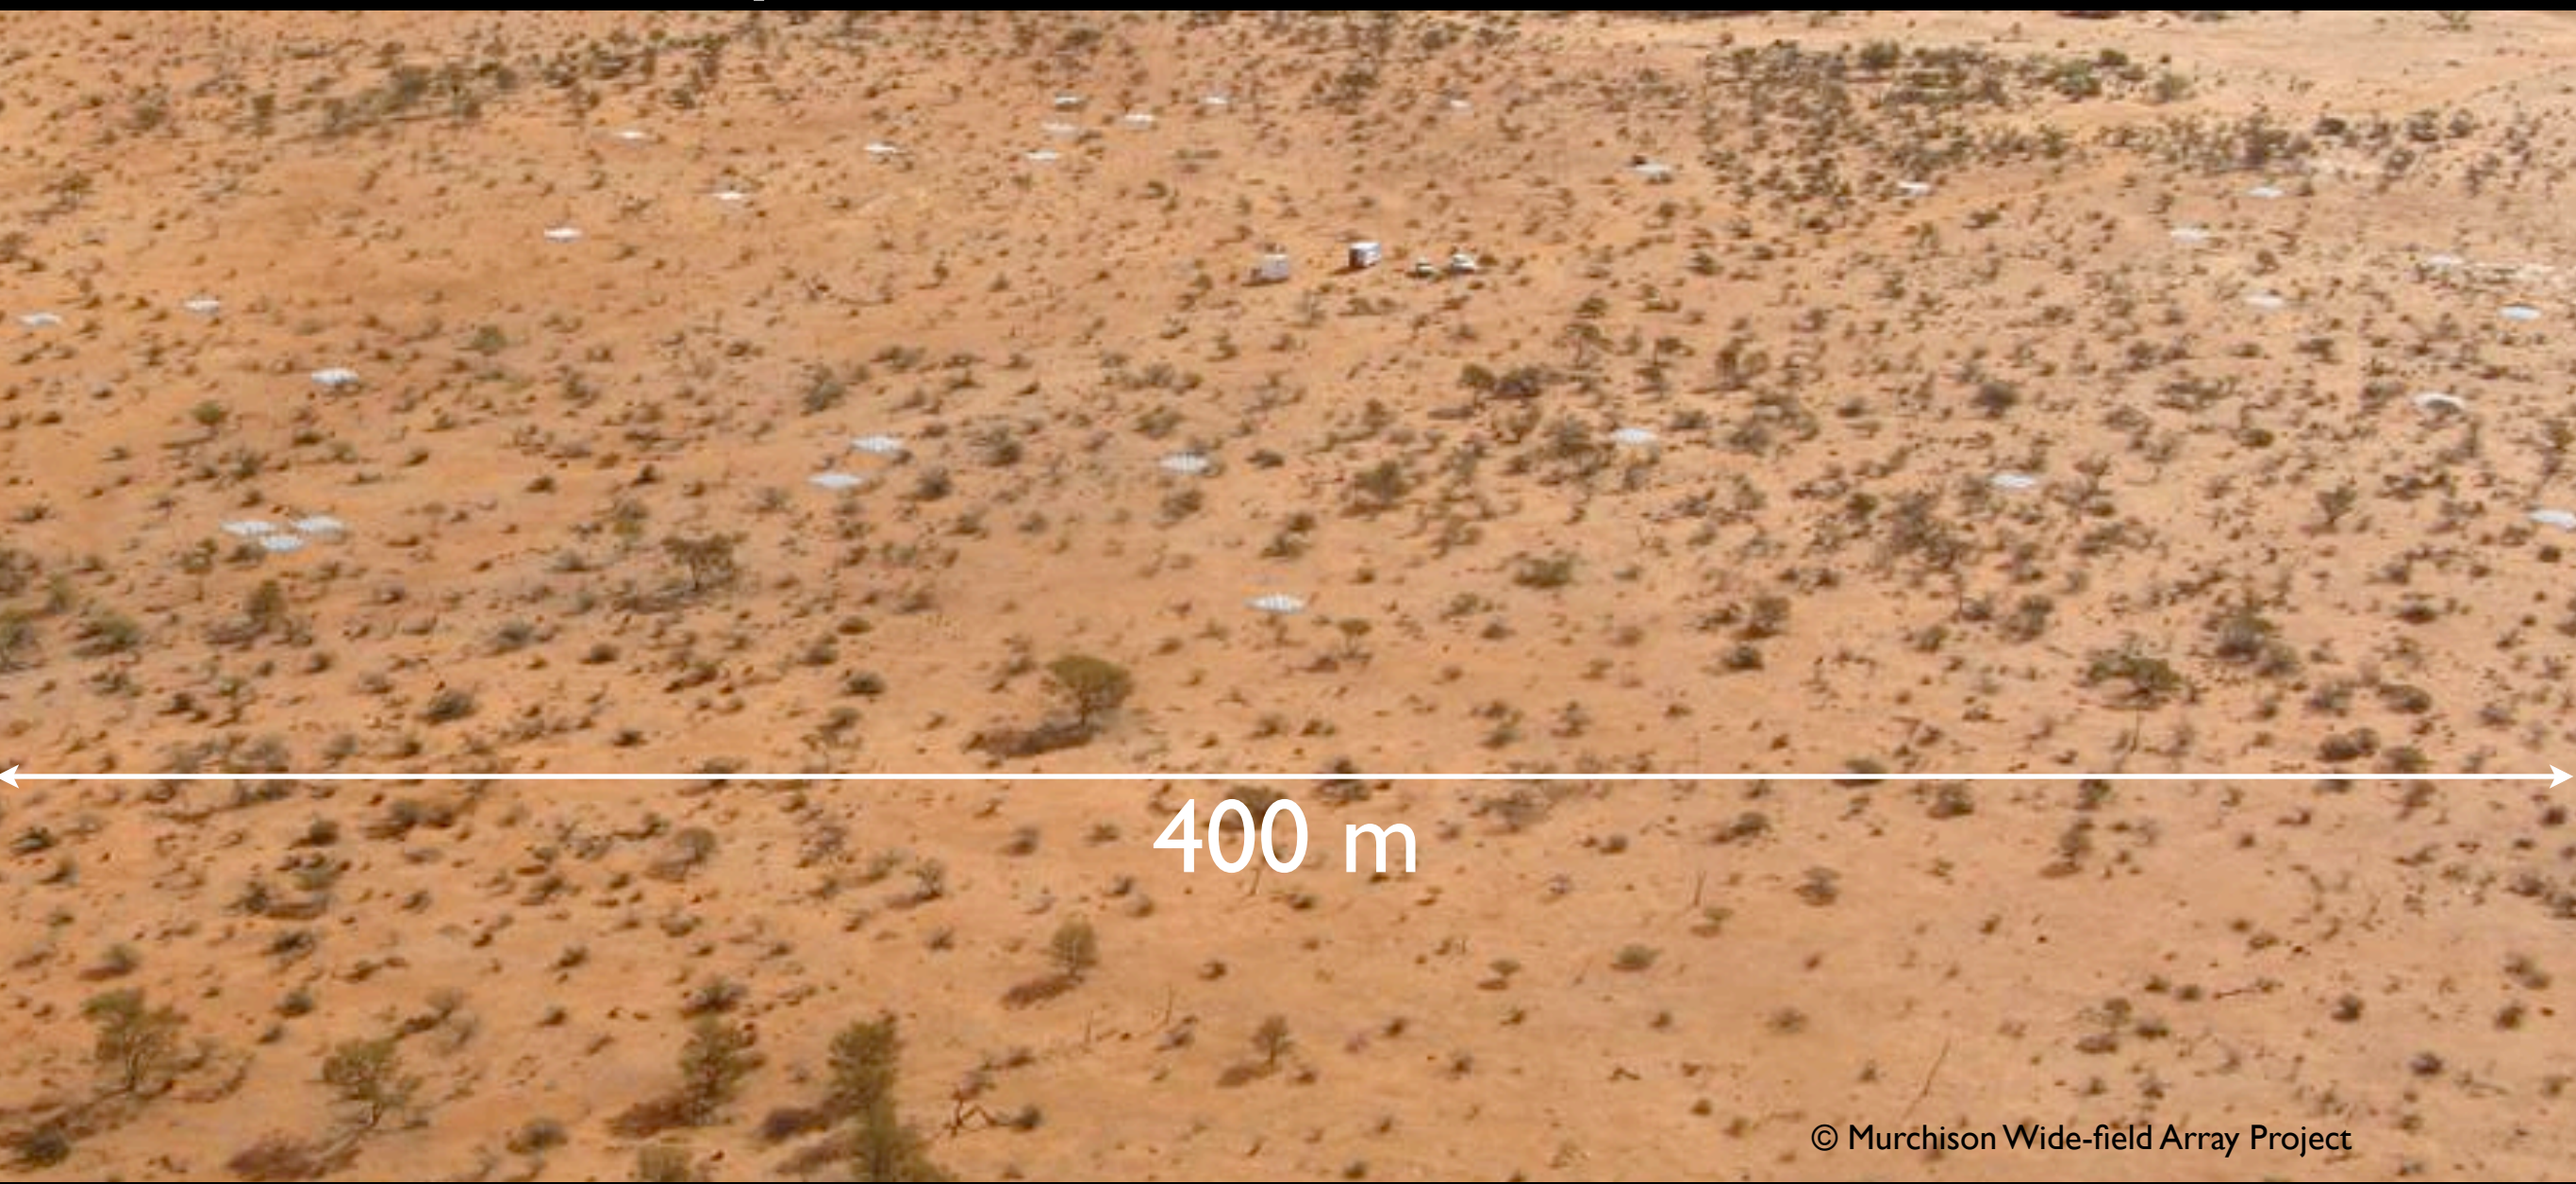

Nowhere (Middle of)
300 km

THANK YOU FOR BEING RADIO QUIET

Outback Computing

- Max. ambient on site: 50° C / 122° F

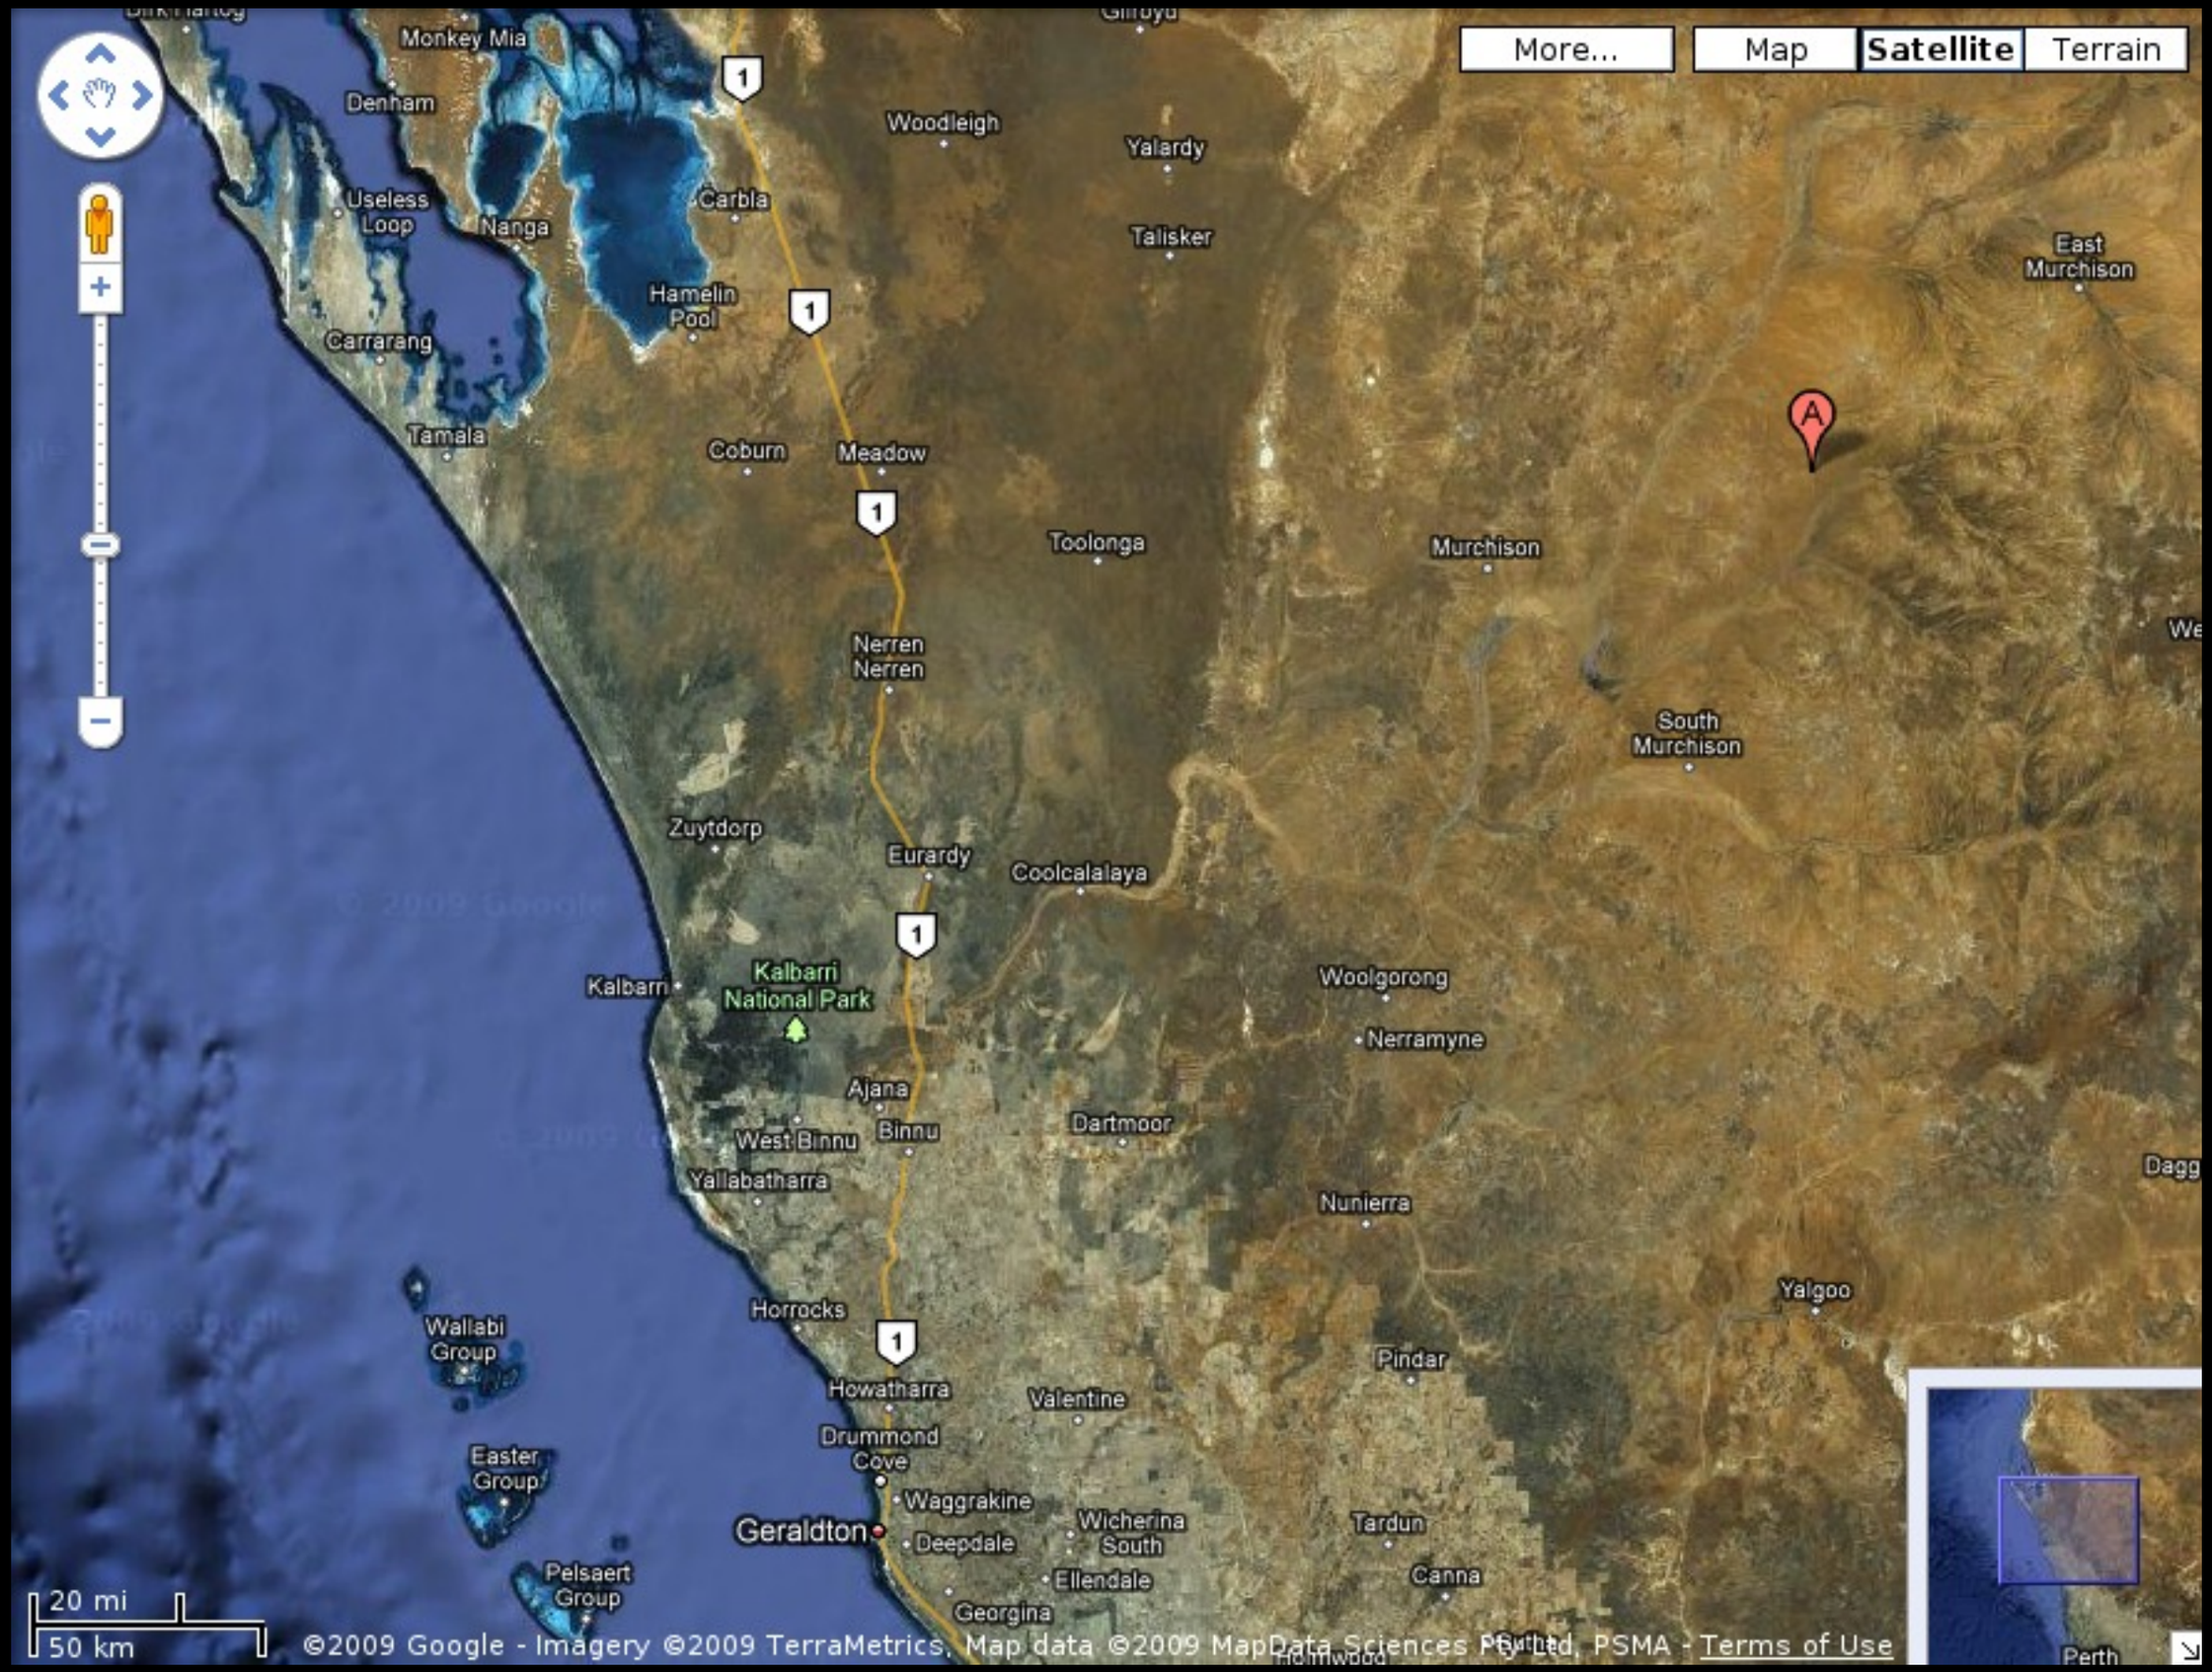
More...
Map
Satellite
Terrain
Monkey Mia
Denham
Useless Loop
Nanga
Carbla
Hamelin Pool
Carrarang
Tamala
Coburn
Meadow
Woodleigh
Yalardy
Talisker
East Murchison
Toolonga
Murchison
Nerren Nerren
South Murchison
Zuytdorp
Eurardy
Coolcalalaya
Kalbarri
Kalbarri National Park
Woolgorong
Nerramyne
Ajana
West Binnu
Binnu
Dartmoor
Yallabatharra
Nunierra
Yalgoo
Horrocks
Wallabi Group
Howatharra
Valentine
Pindar
Drummond Cove
Easter Group
Waggrakine
Geraldton
Deepdale
Wicherina South
Tardun
Ellendale
Canna
Pelsaert Group
Georgina
20 mi
50 km
©2009 Google - Imagery ©2009 TerraMetrics, Map data ©2009 MapData Sciences Pty Ltd, PSMA - Terms of Use
Perth

N

324 **Mullewa Carnarvon Rd**, Murchison, WA, Australia
Address is approximate

Mullewa Carnarvon Rd

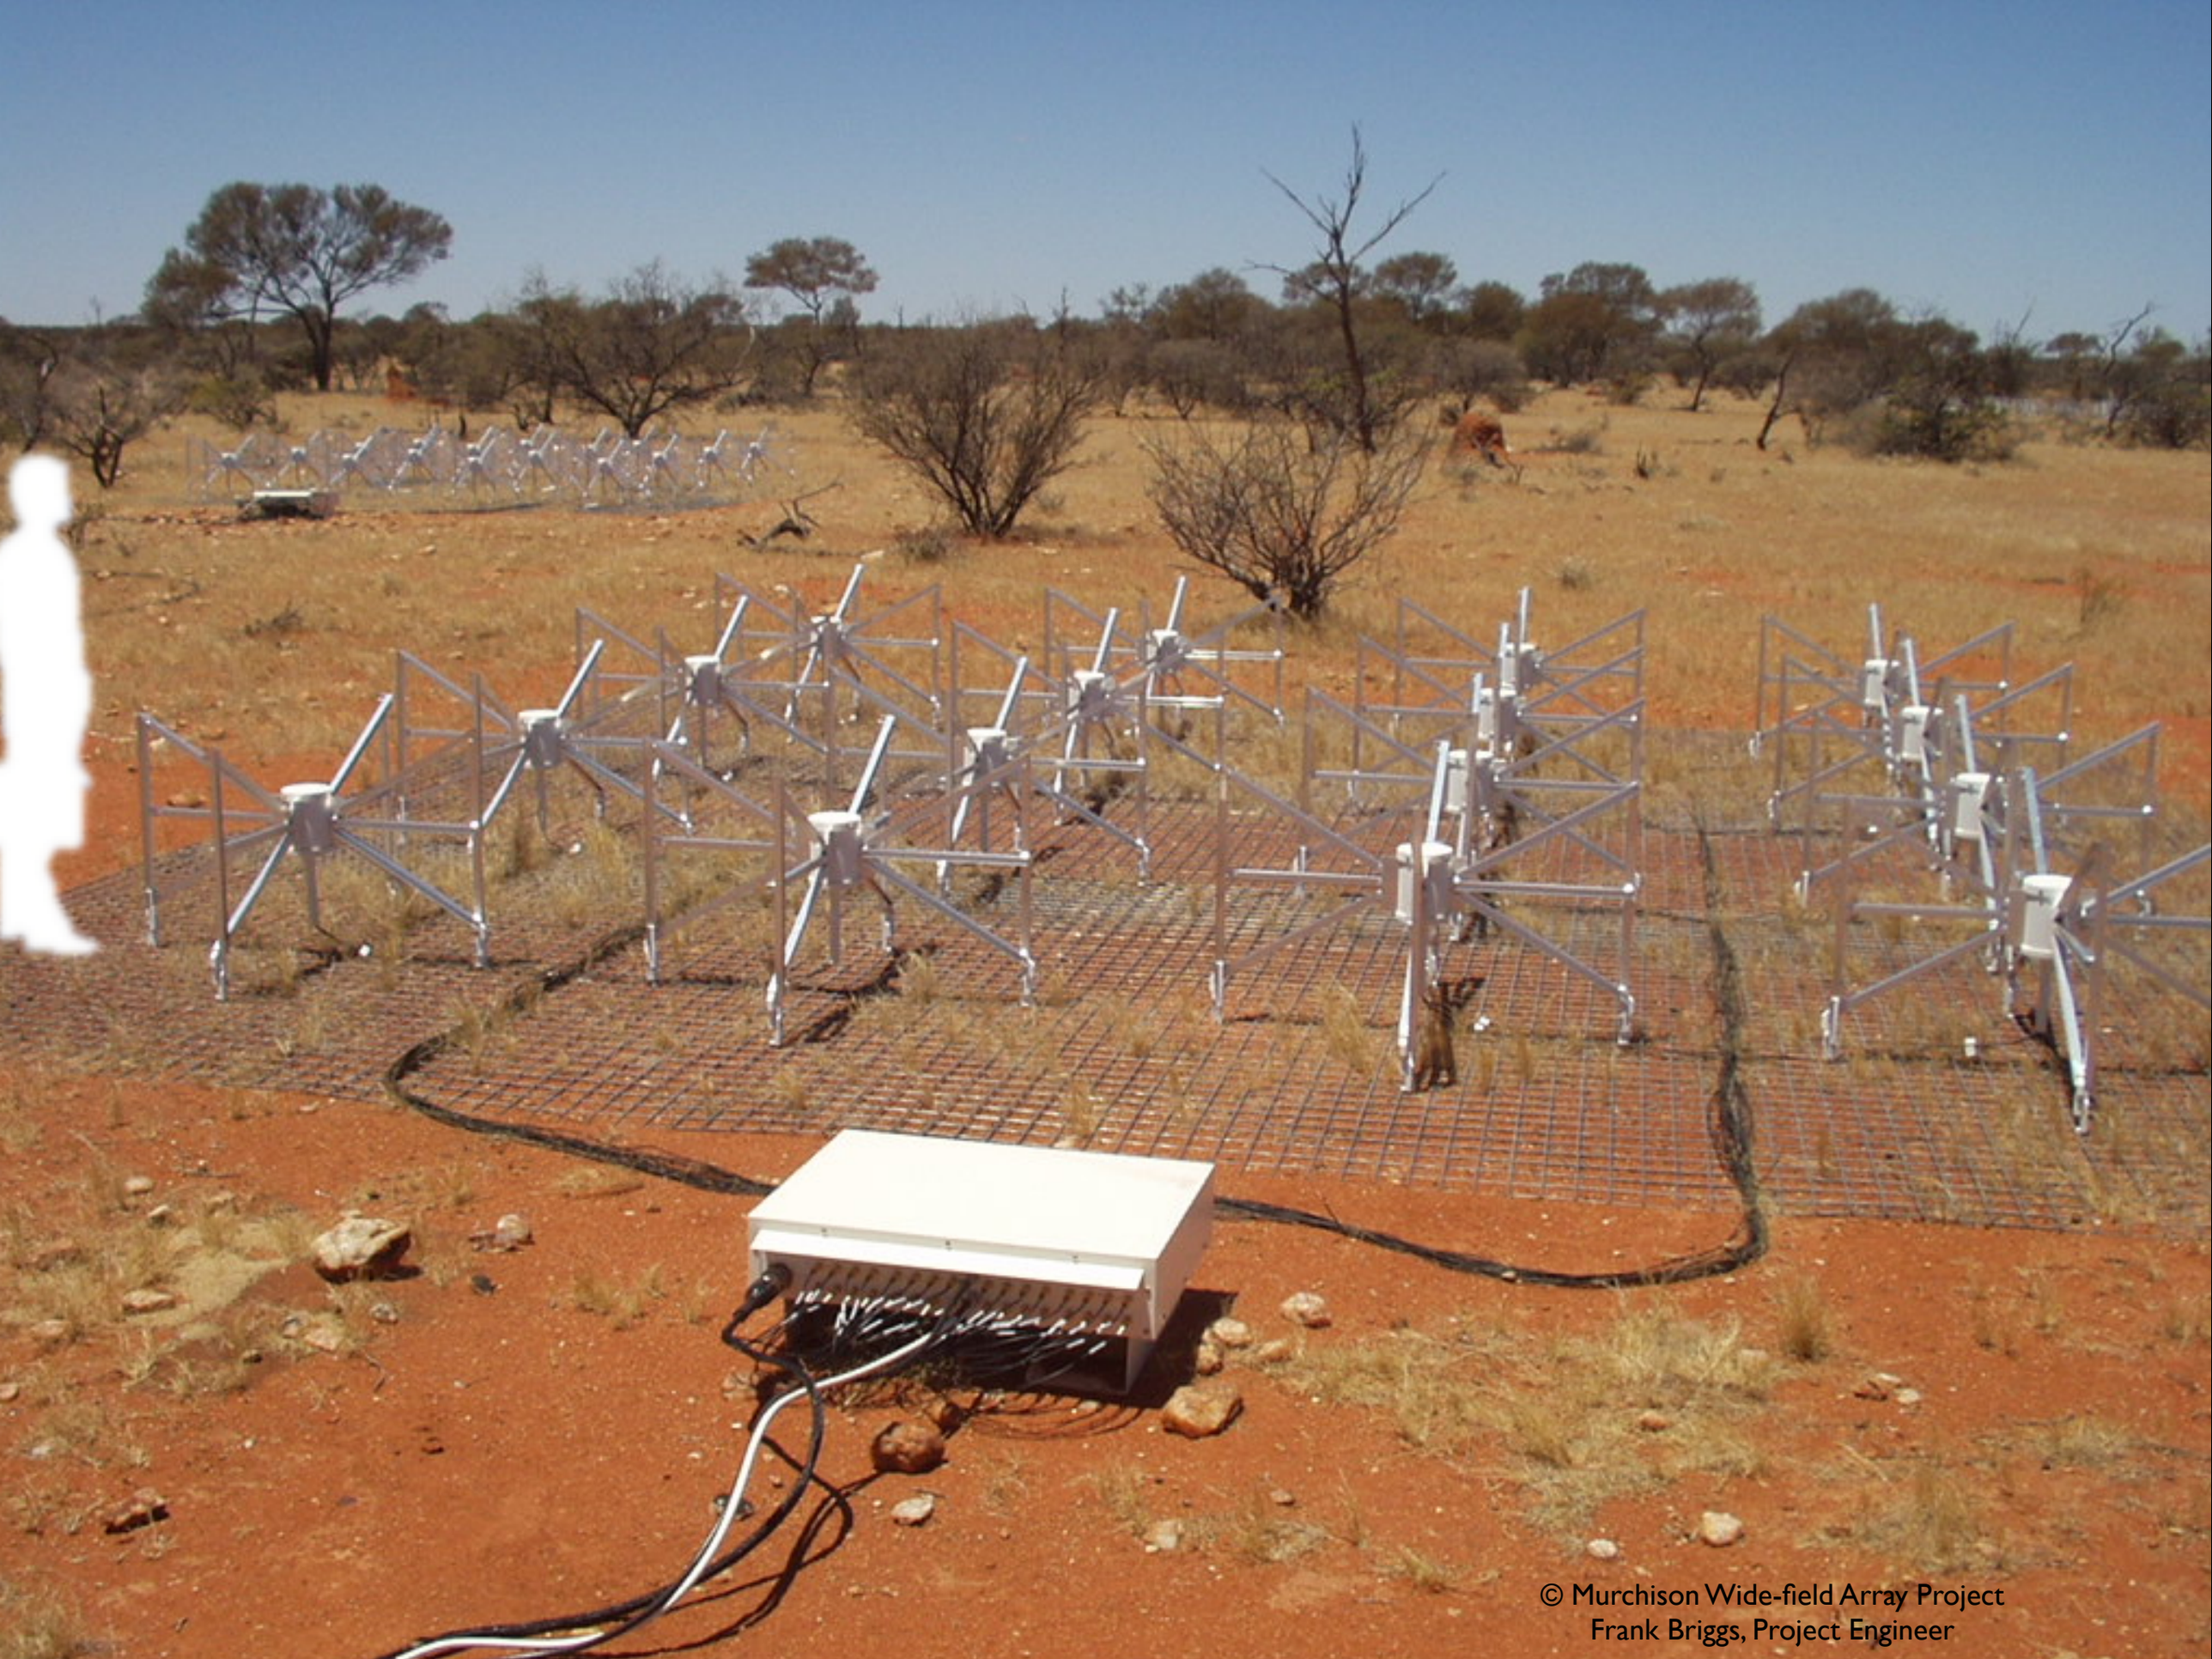

© Murchison Wide-field Array Project
Frank Briggs, Project Engineer

Nervous System

Supercomputing Center

Who is the GPU expert?

Who is the GPU expert?

Life Support

MWA Pipeline

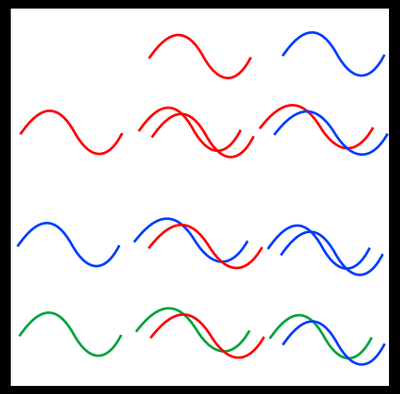

160 Gb s^{-1}

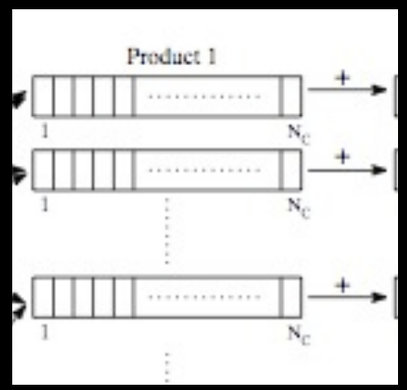

5Gb s^{-1}

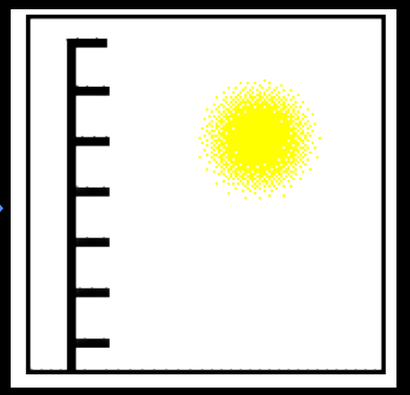

Correlator

20 TFLOP @ 20 kW

5Gb s^{-1}

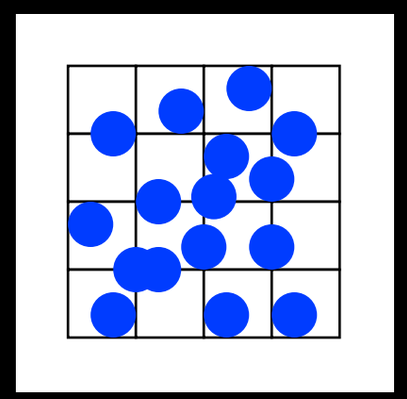

8 Gb s^{-1}

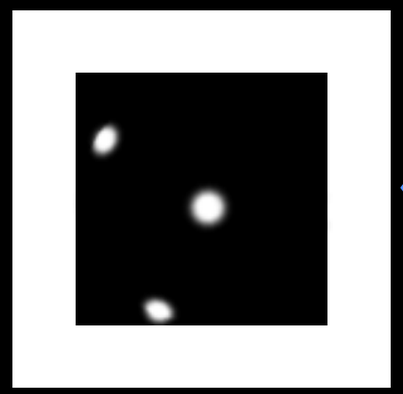

4 Gb s^{-1}

Image

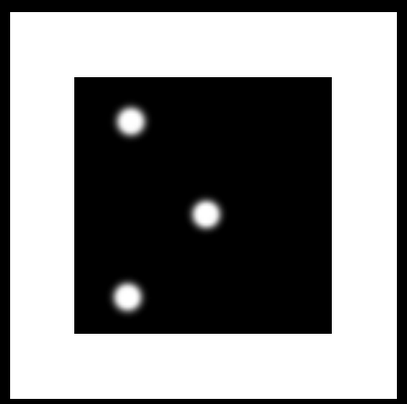

20 Degrees FOV

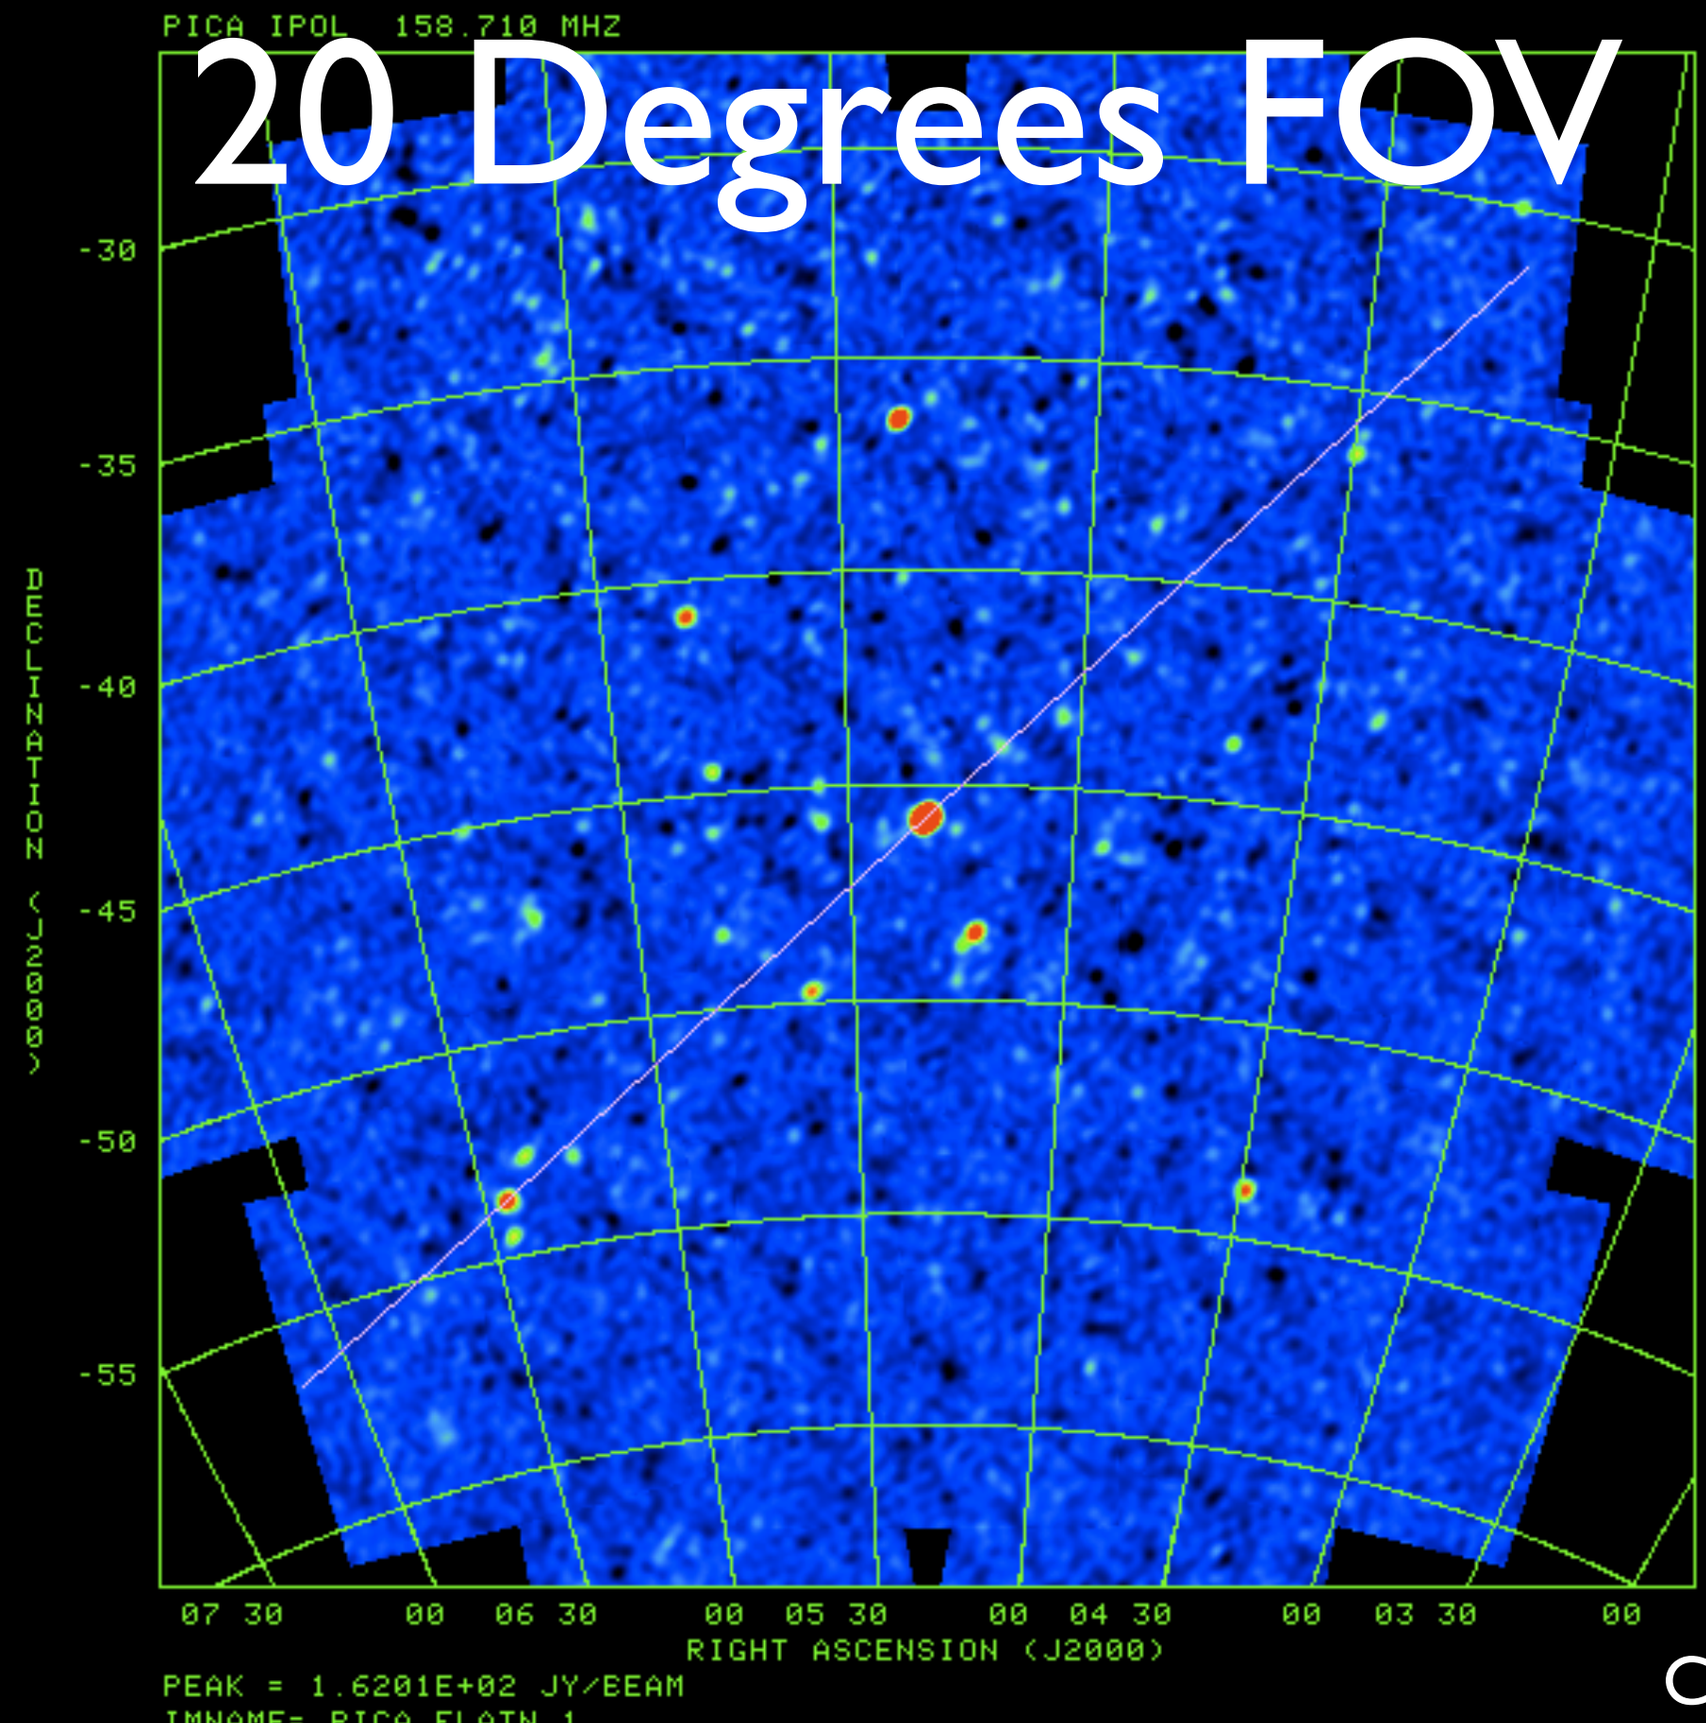

C. Williams, MIT

Real Time System

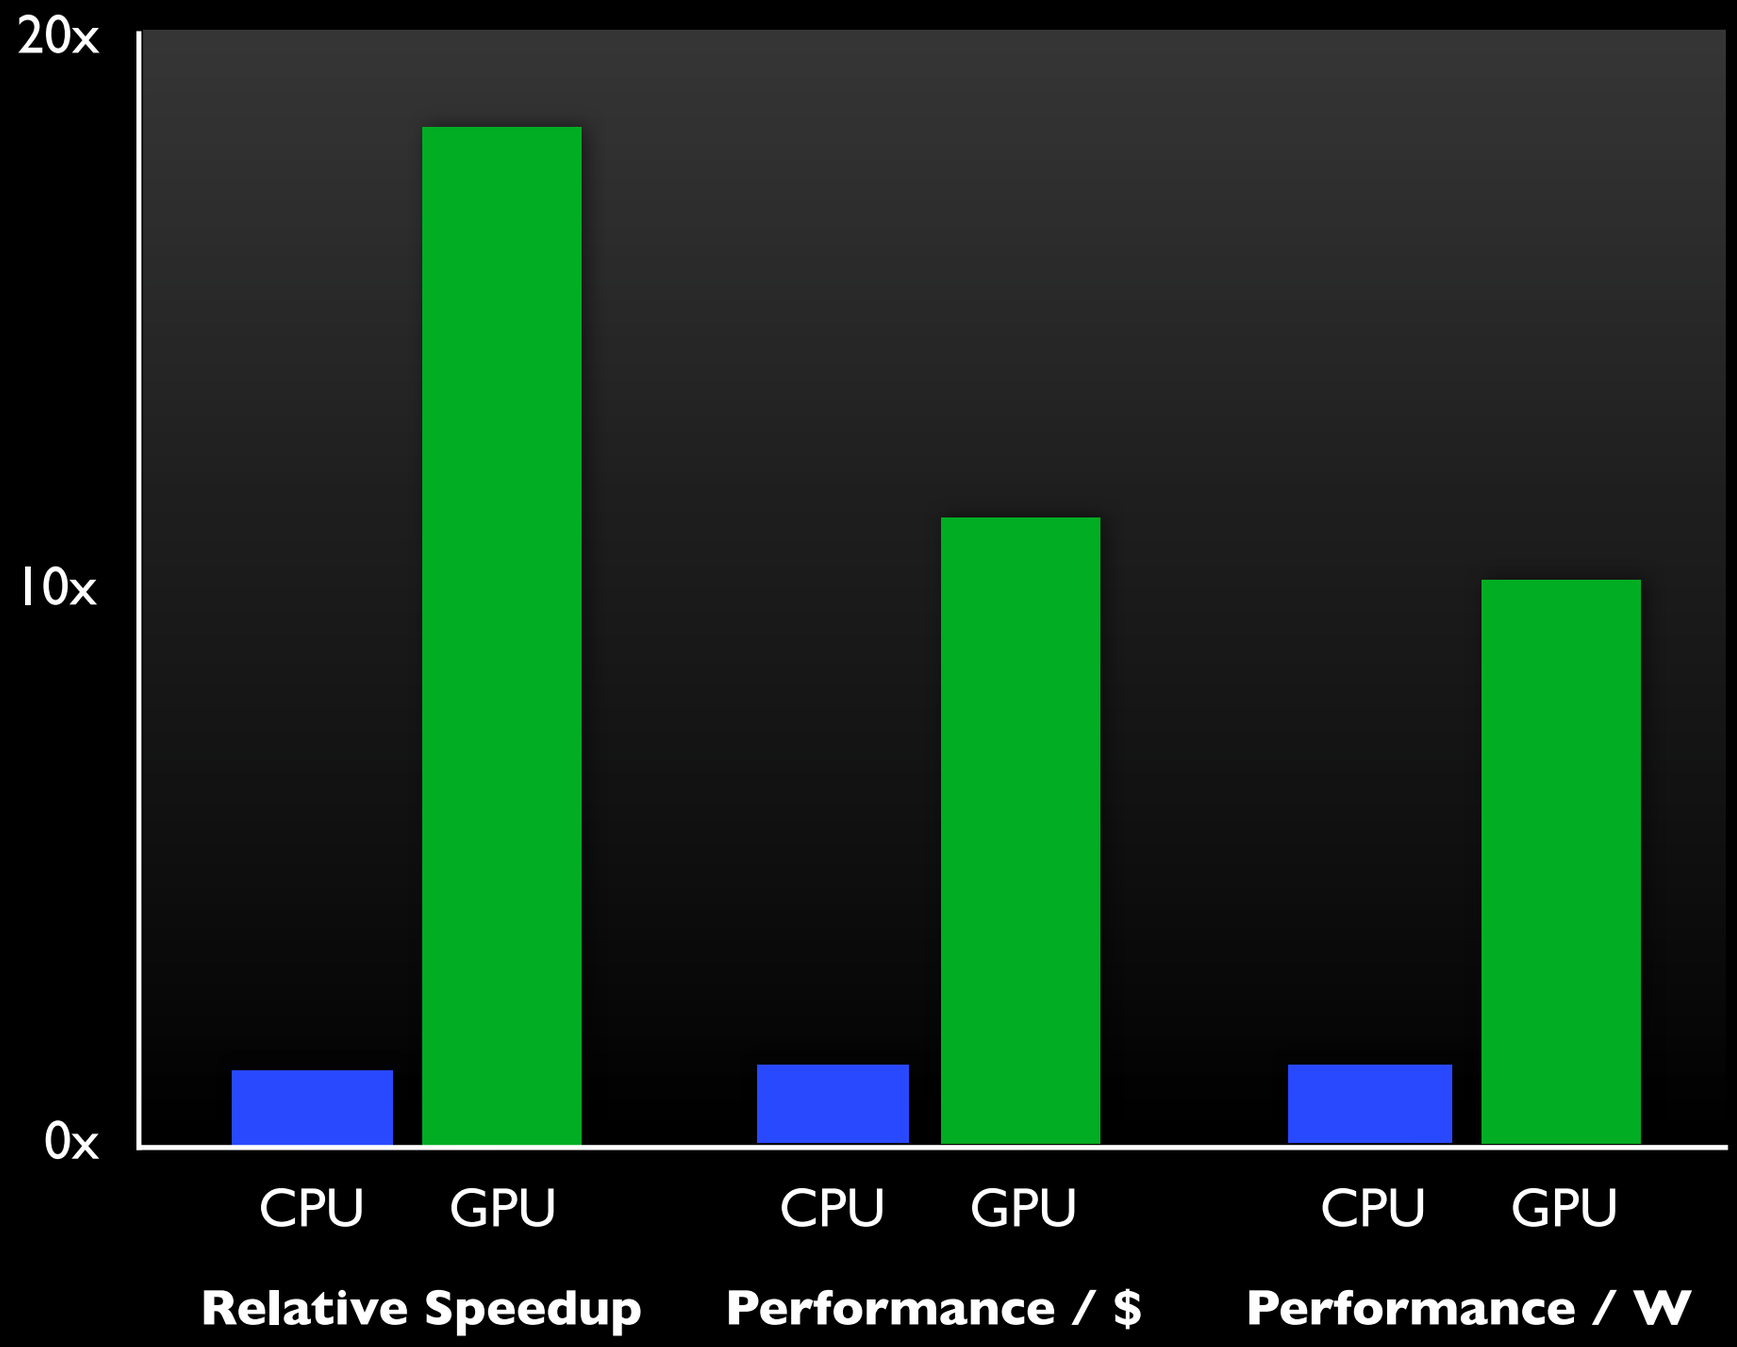

High-Throughput Radio Astronomy

Steer Antenna

010011....

Images

Antenna Array

Interferometry

Science

Real-Time Steering

- Predict solar flares and coronal mass ejections

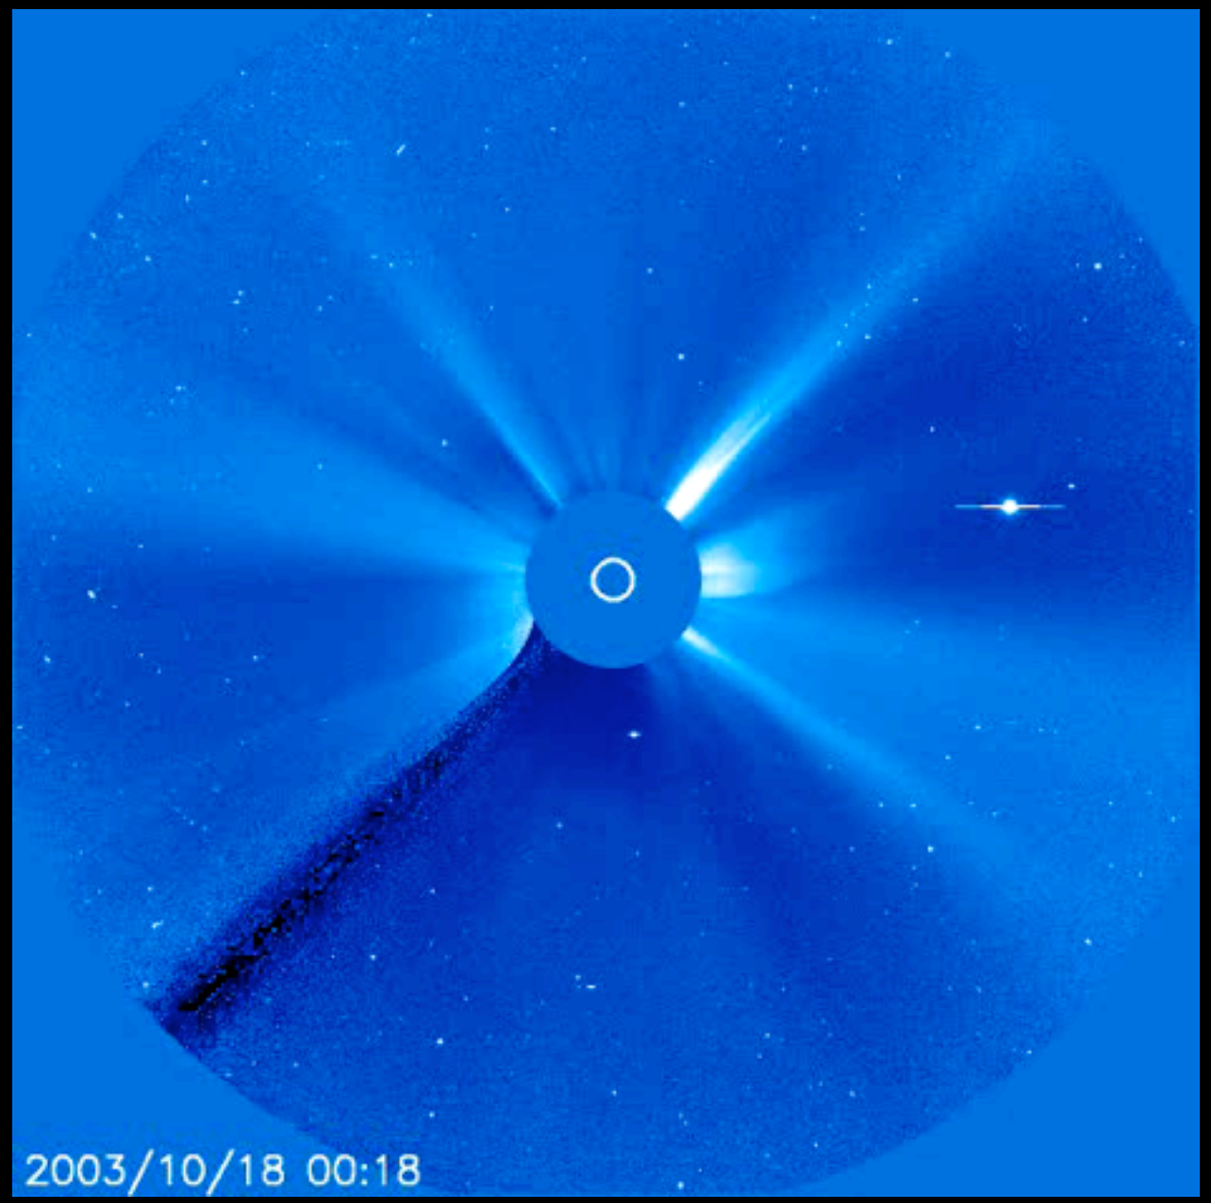

SOHO/EIT (ESA & NASA)

The Future

Square Kilometer Array

1 EFLOP / 1 Tb s^{-1} ca. 2020

More Information

Richard Edgar
Diesel-Powered GPU Computing: Enabling a Real-Time Radio Telescope in the Australian Outback
Session Id#1065
Friday 3 pm California

Simulations

From quantum chemistry to physics

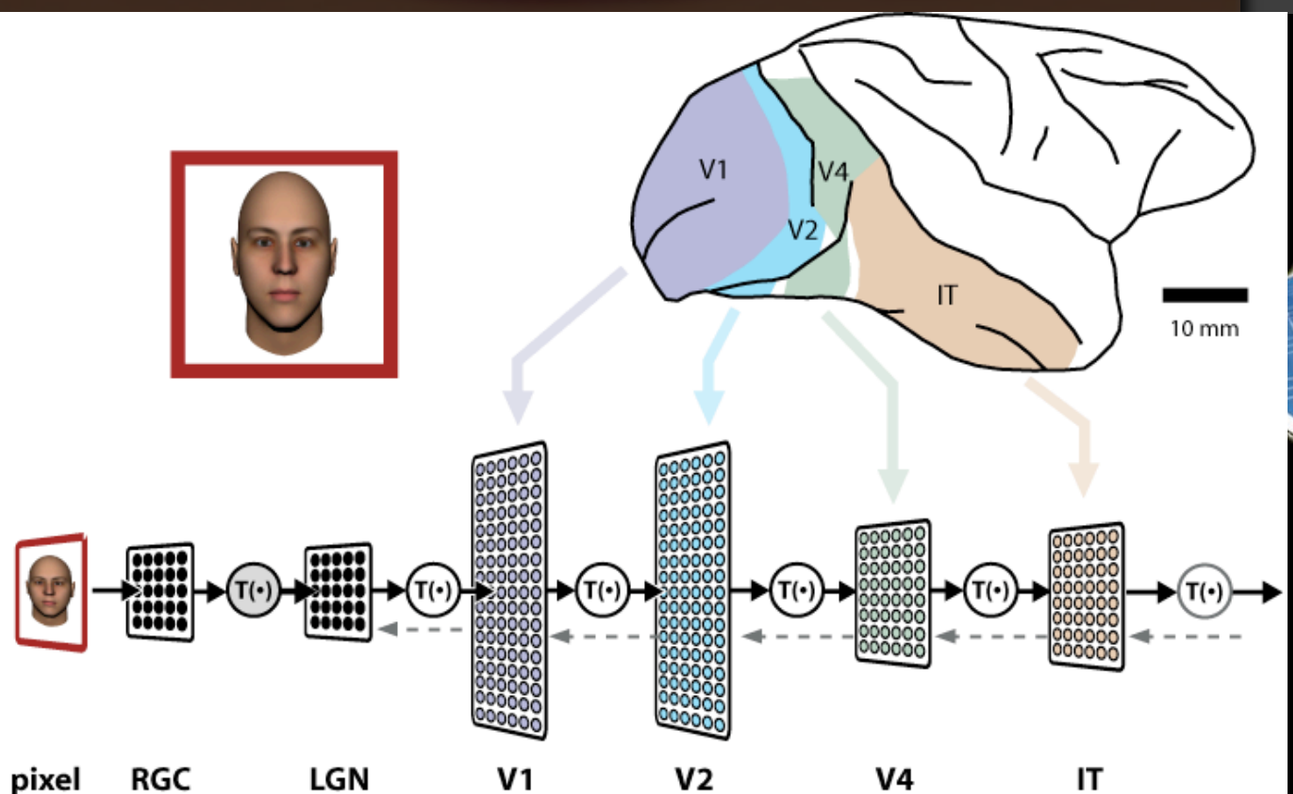

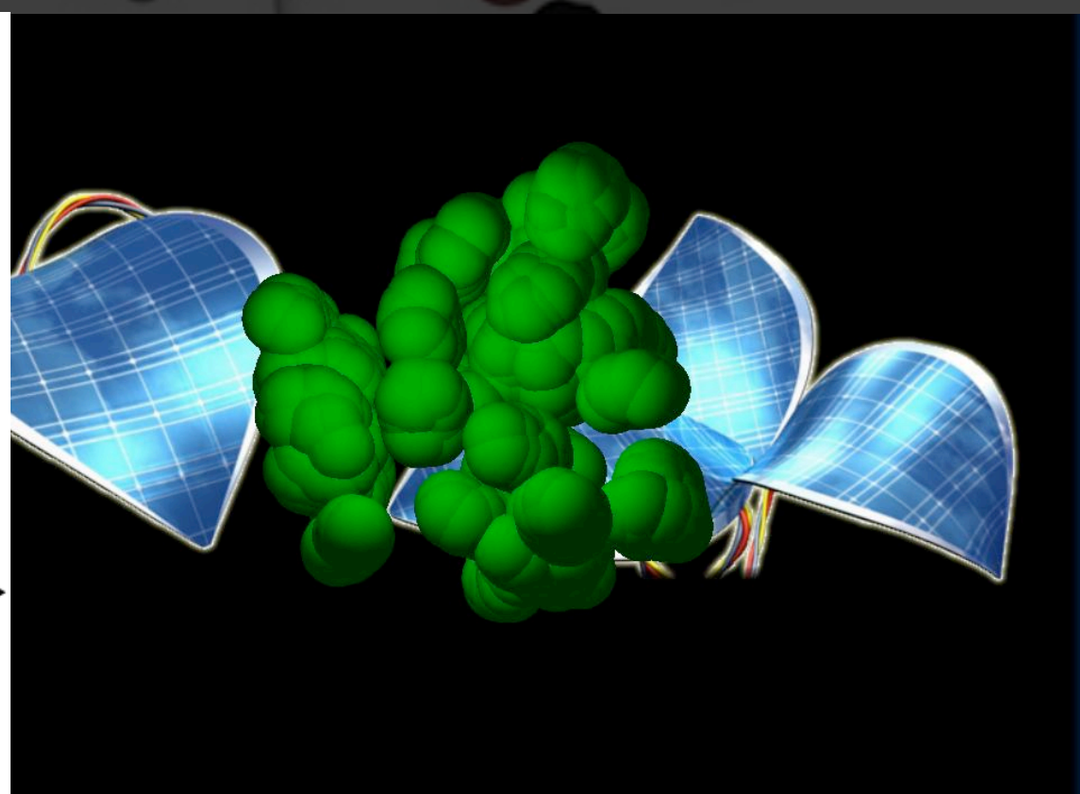

Quantum Chemistry

- Alan Aspuru-Guzik
- Mark Watson
- Richard Edgar
- Kenta Hongo
- Roberto Olivares
- Leslie Vogt
- Sean Kermes


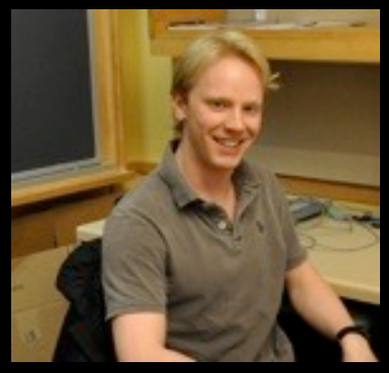





The Scientific Challenge

- Determine properties of organic molecules

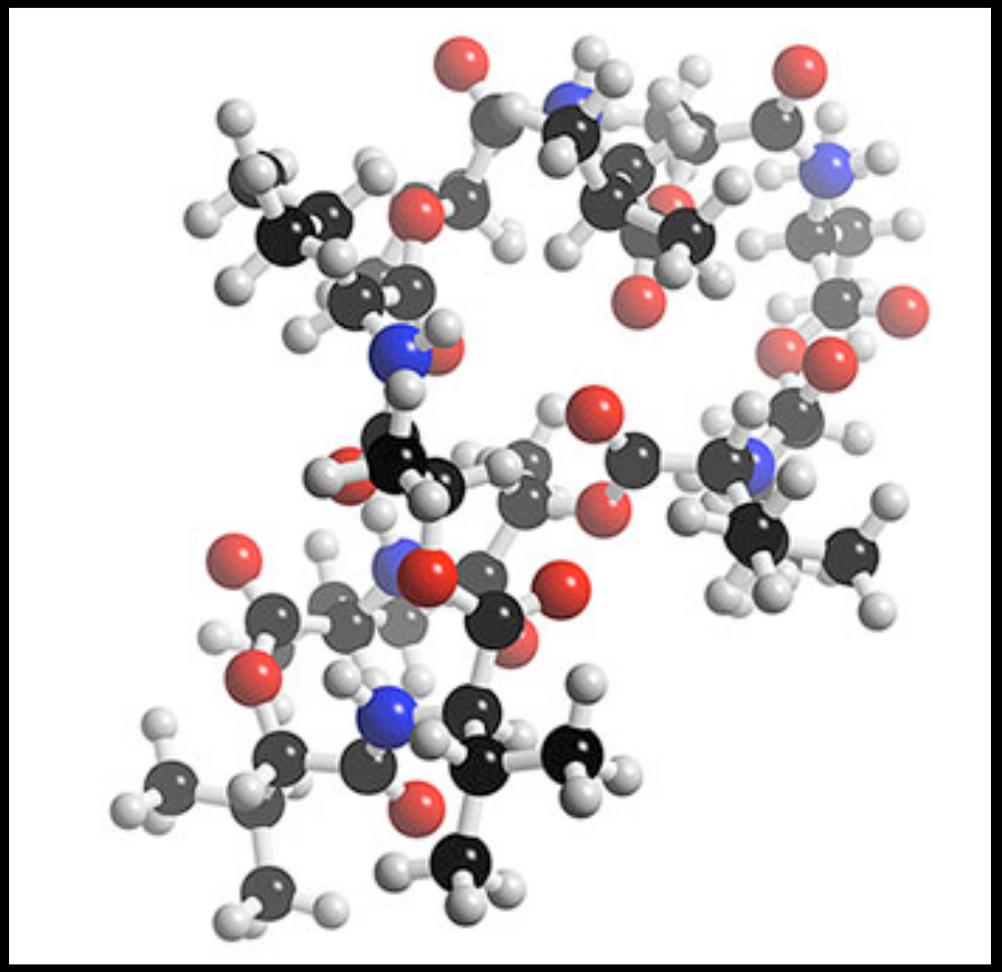

Quantum Many-Body Problems

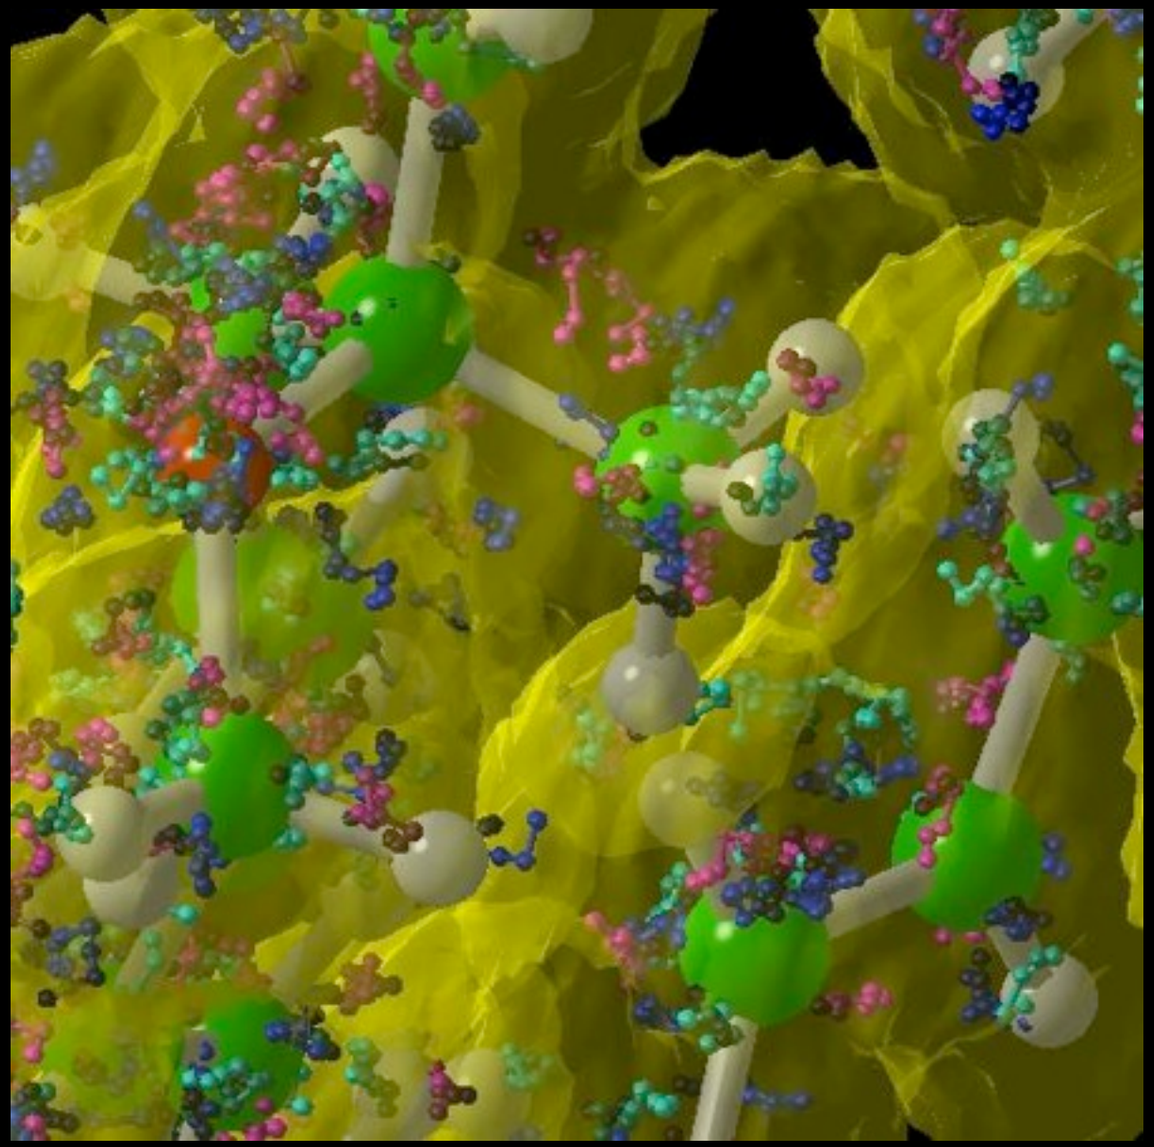

Results

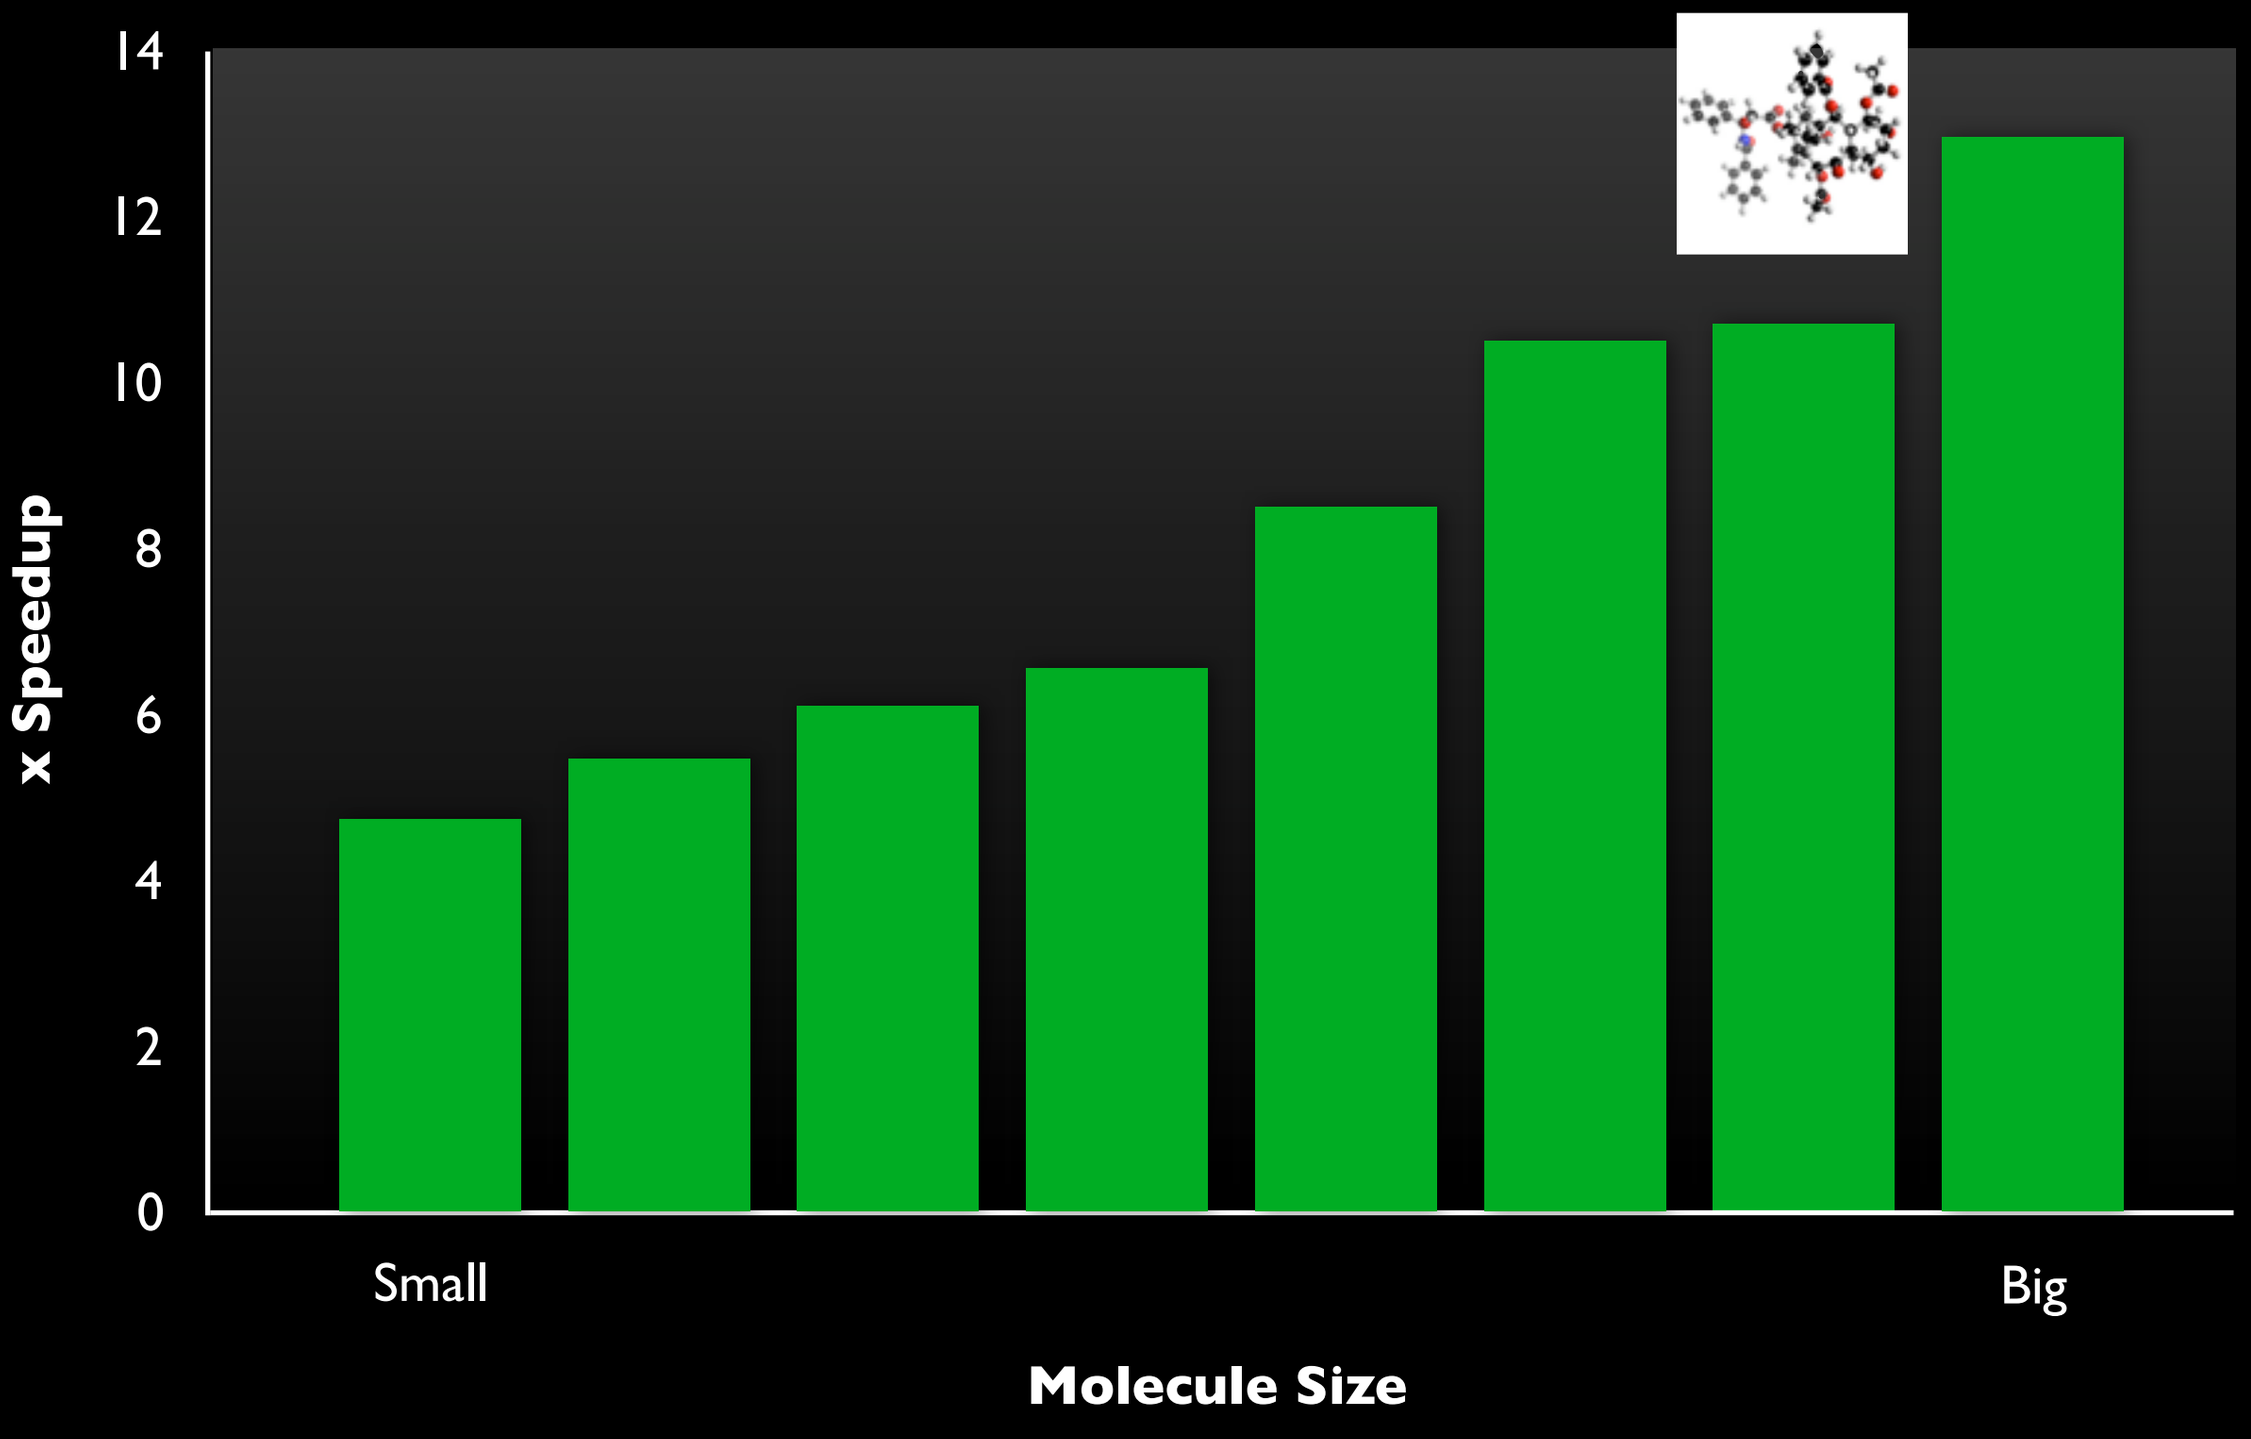

The Future

- Improve organic photovoltaic materials

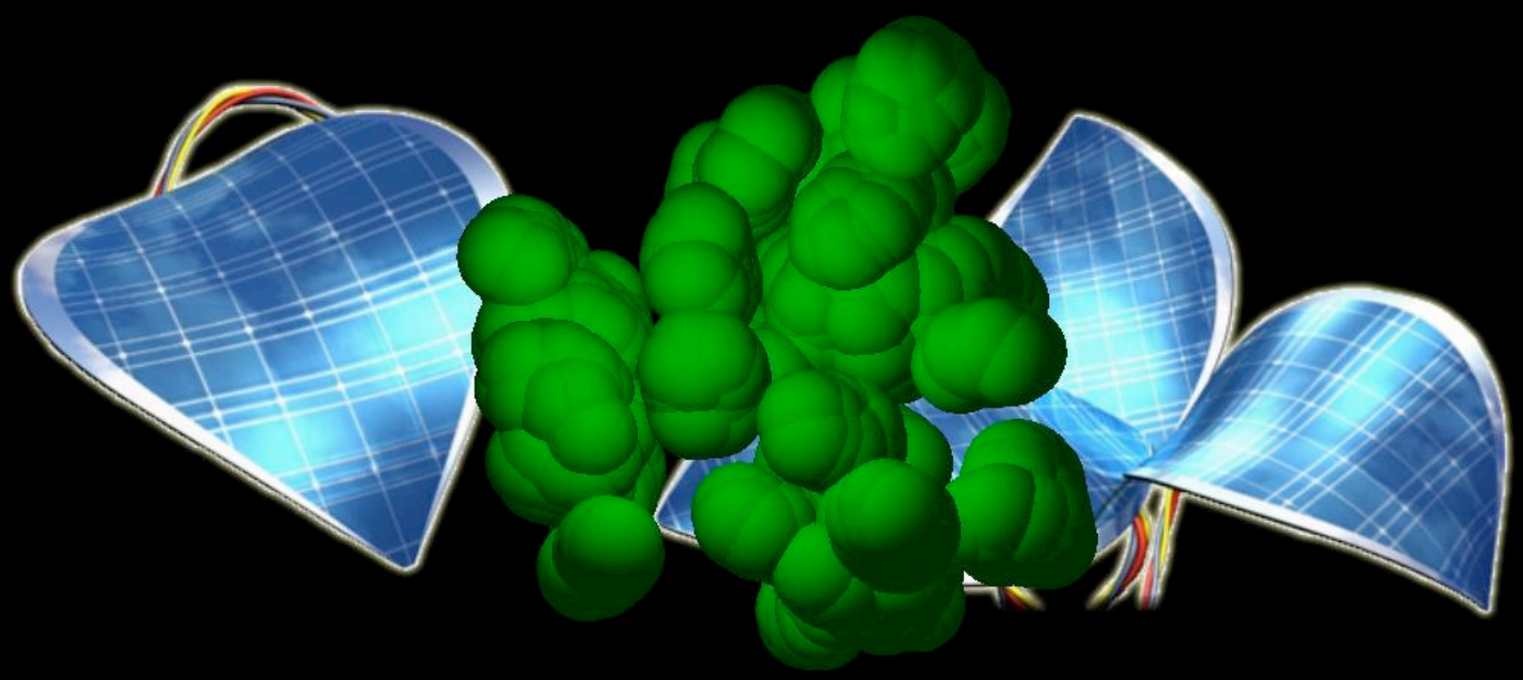

Volunteer Computing

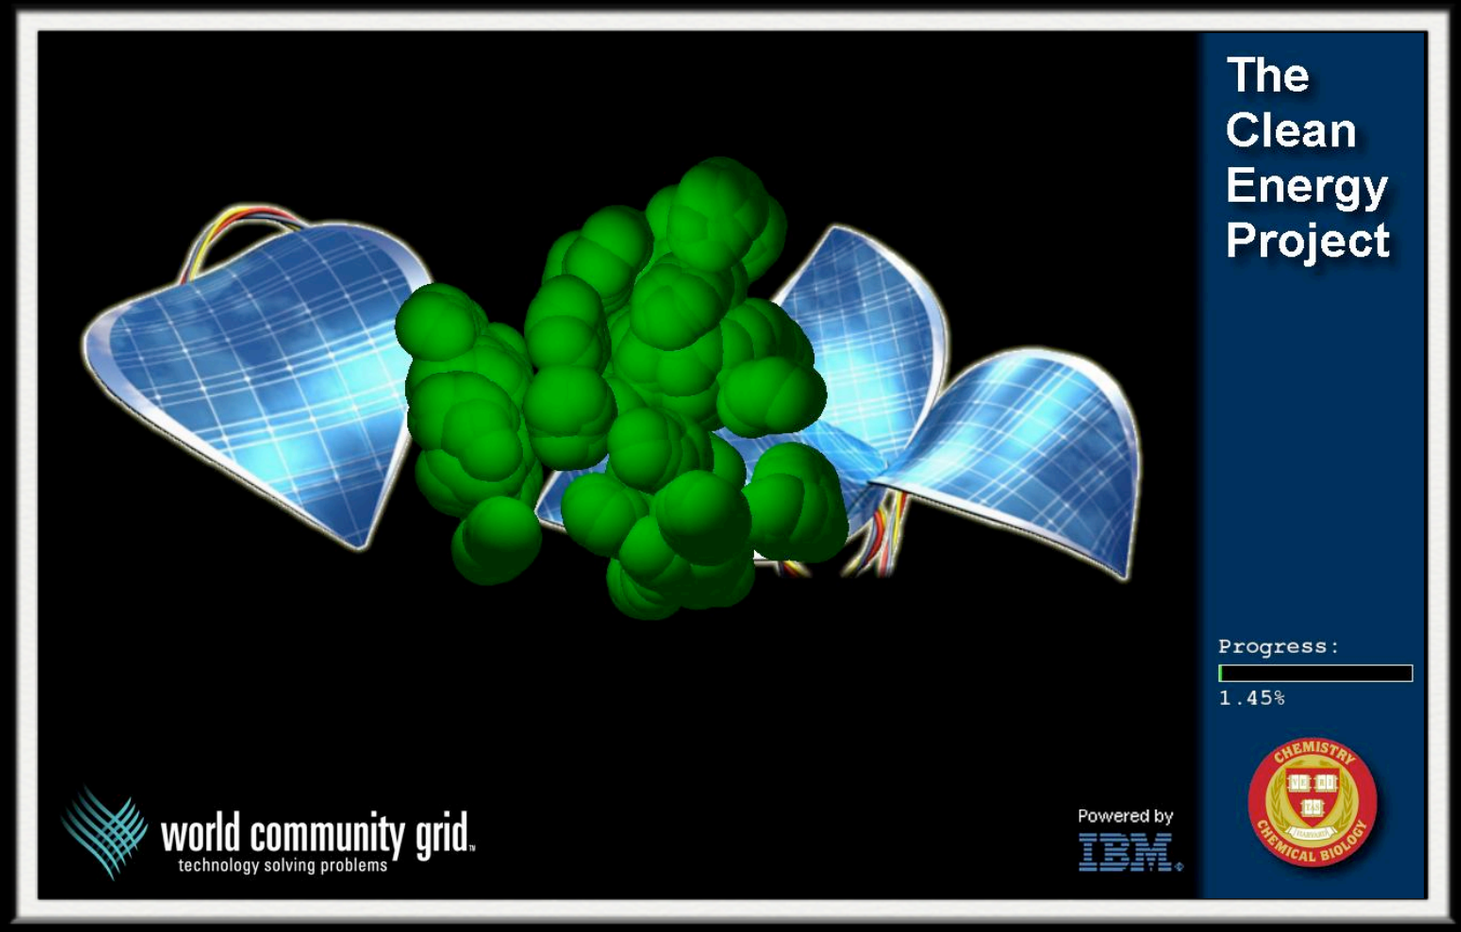

Relevant Talk

David Anderson
Volunteer Computing for GPUs: Petaflops for Free
Session Id#1010
Thursday 3 pm California

Visual Neuroscience

- David Cox, The Rowland Institute at Harvard

- Nicolas Pinto, MIT

- Jim DiCarlo, MIT

The Scientific Challenge

- How does the brain perform object recognition?

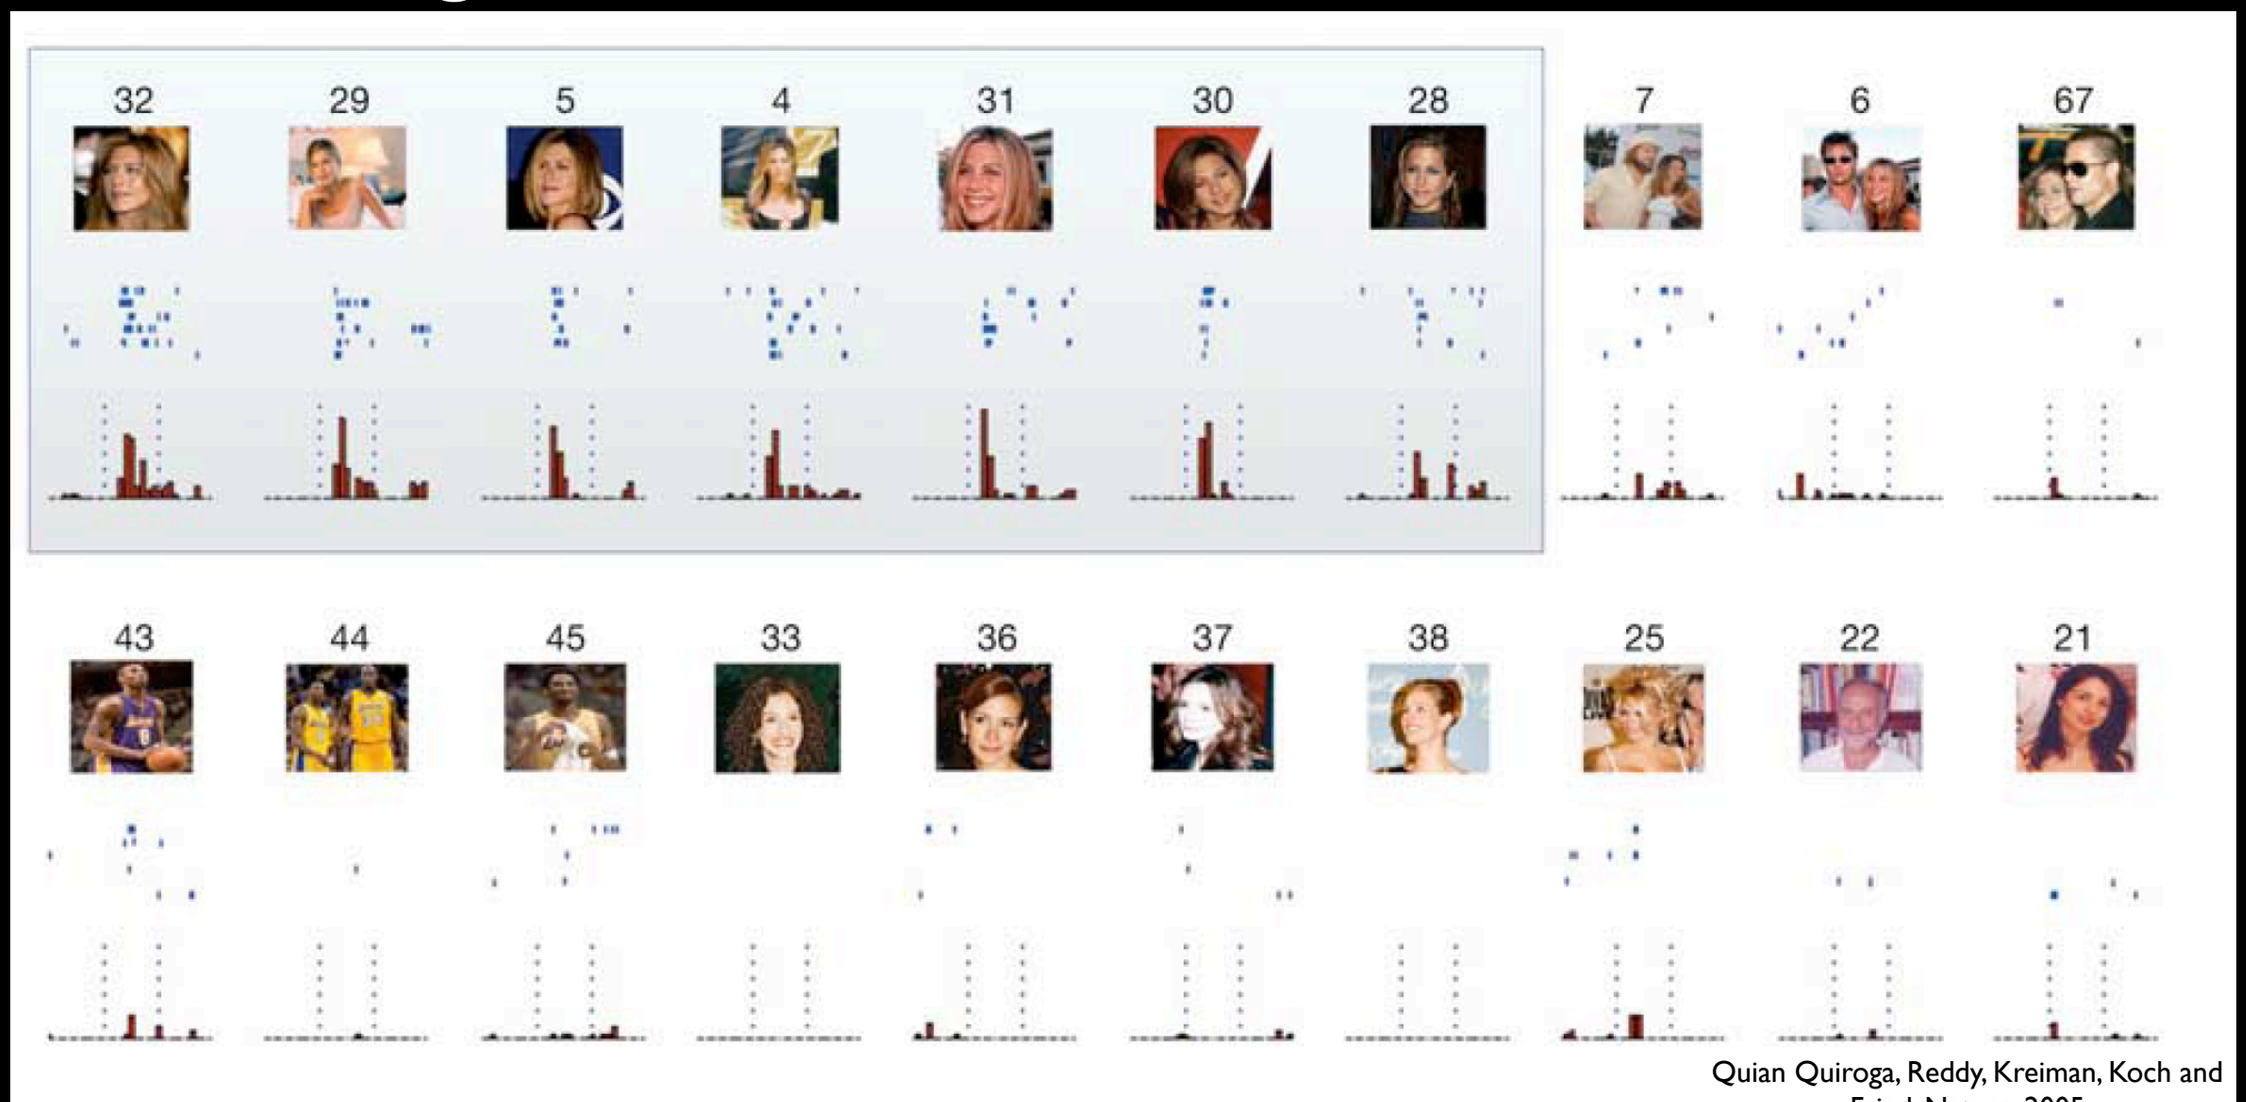

Quian Quiroga, Reddy, Kreiman, Koch and Fried, Nature, 2005

The Scientific Challenge

- What are good models of the human visual system, and what are their parameters?

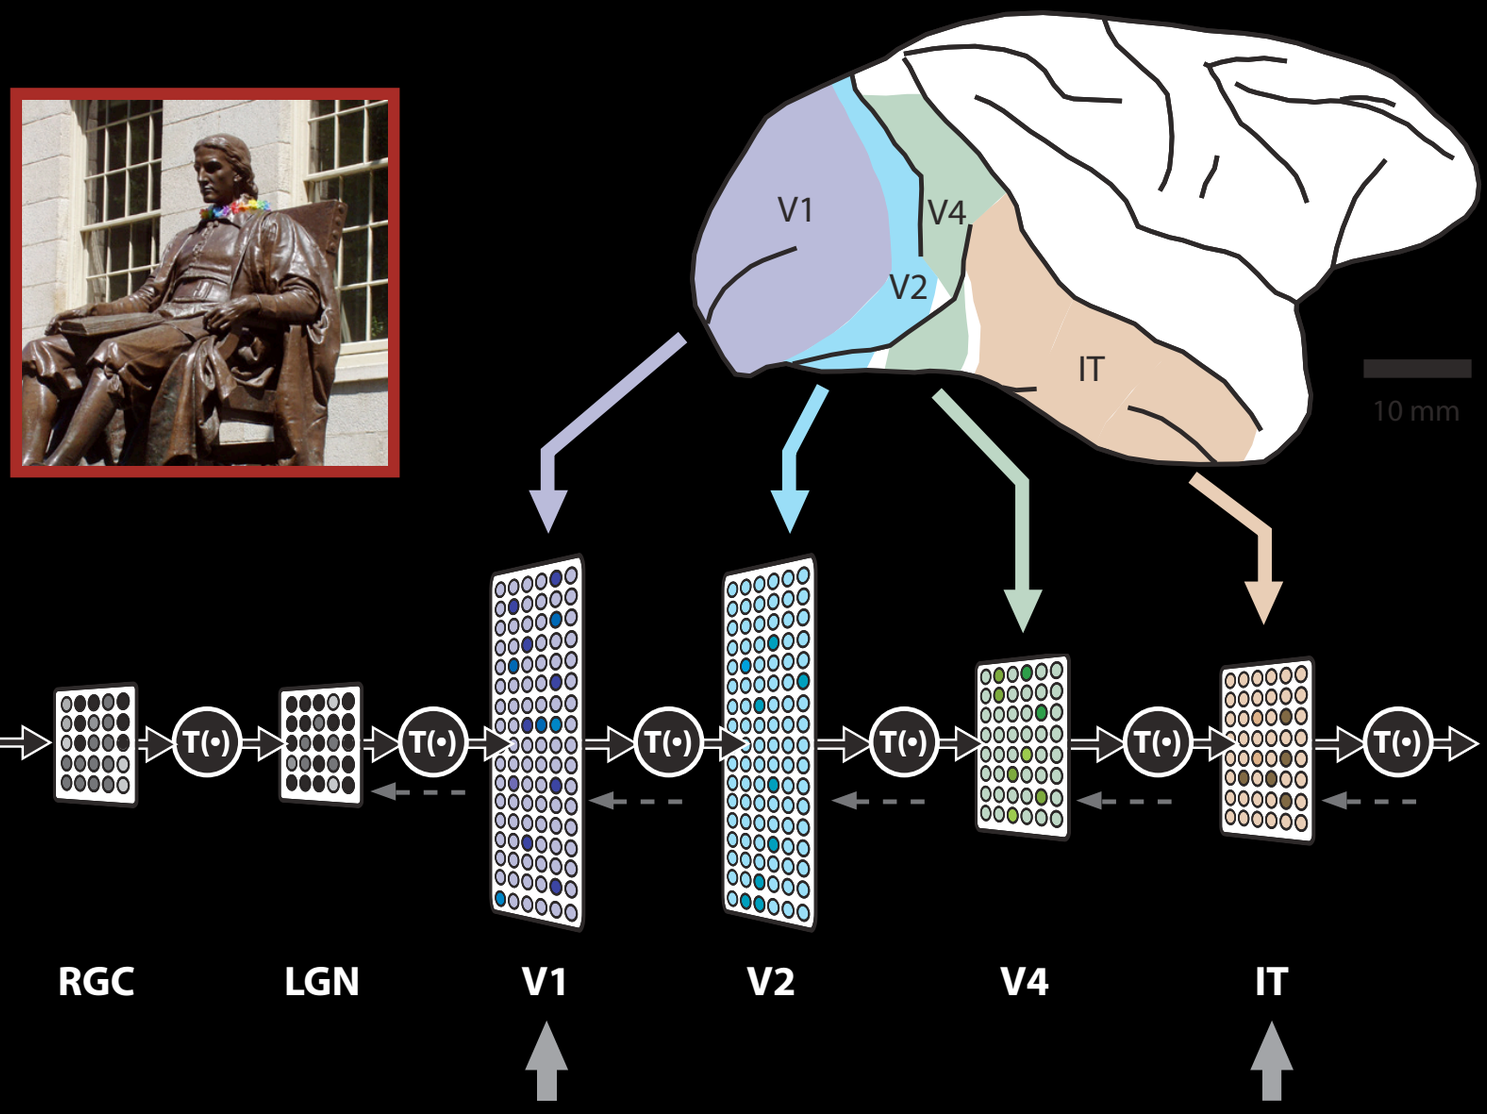

Simulations

Train 10,000 to 100,000 different models
on real image & video data

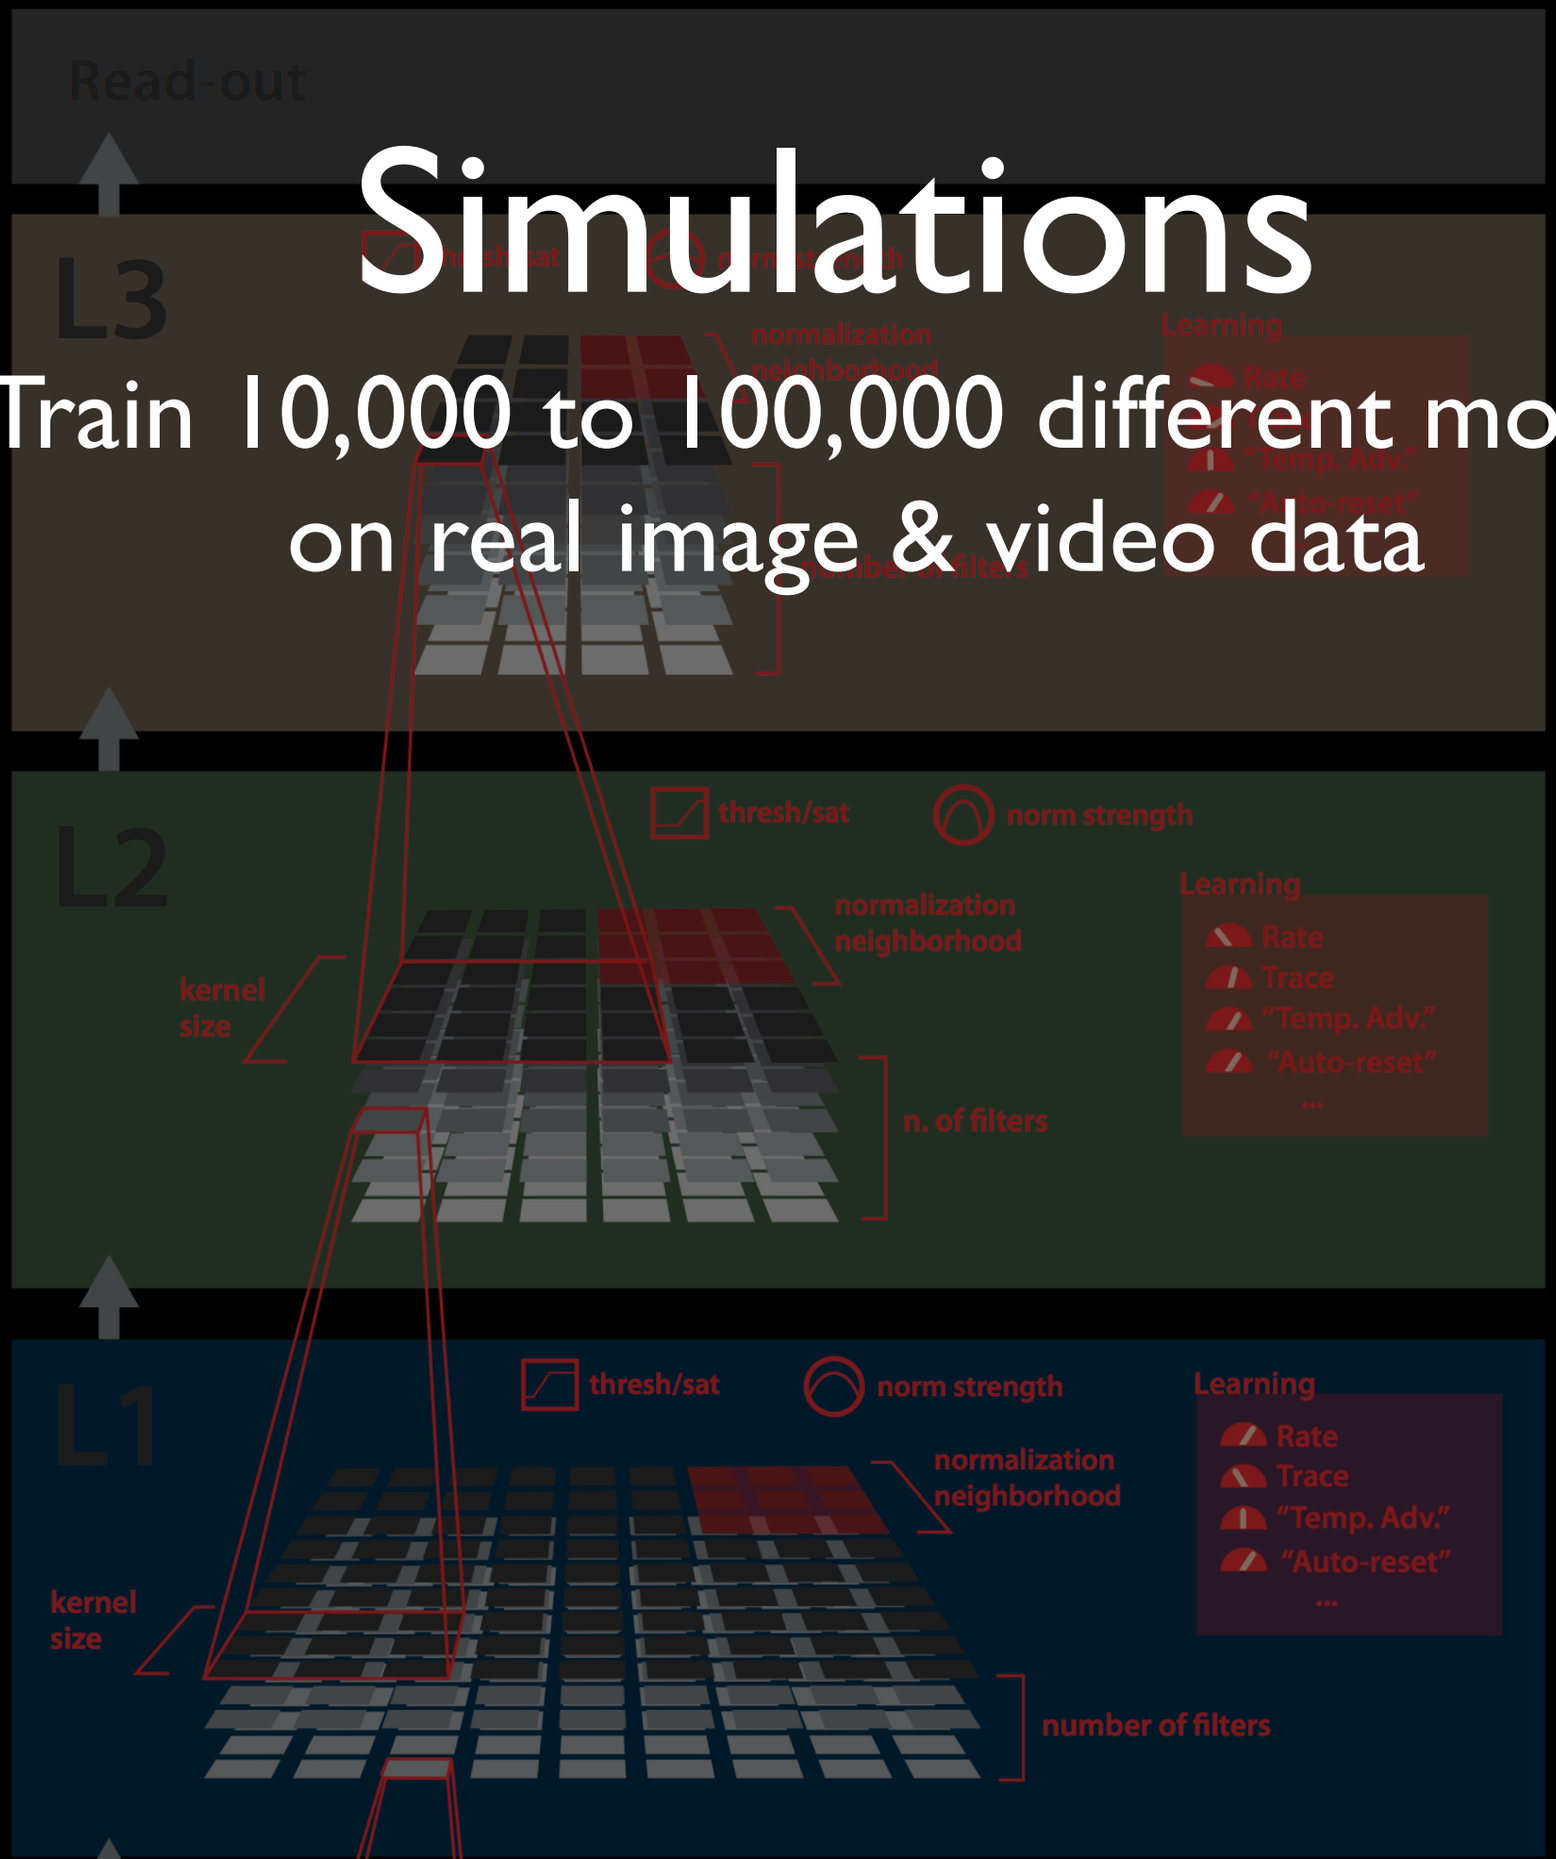

Performance

vs

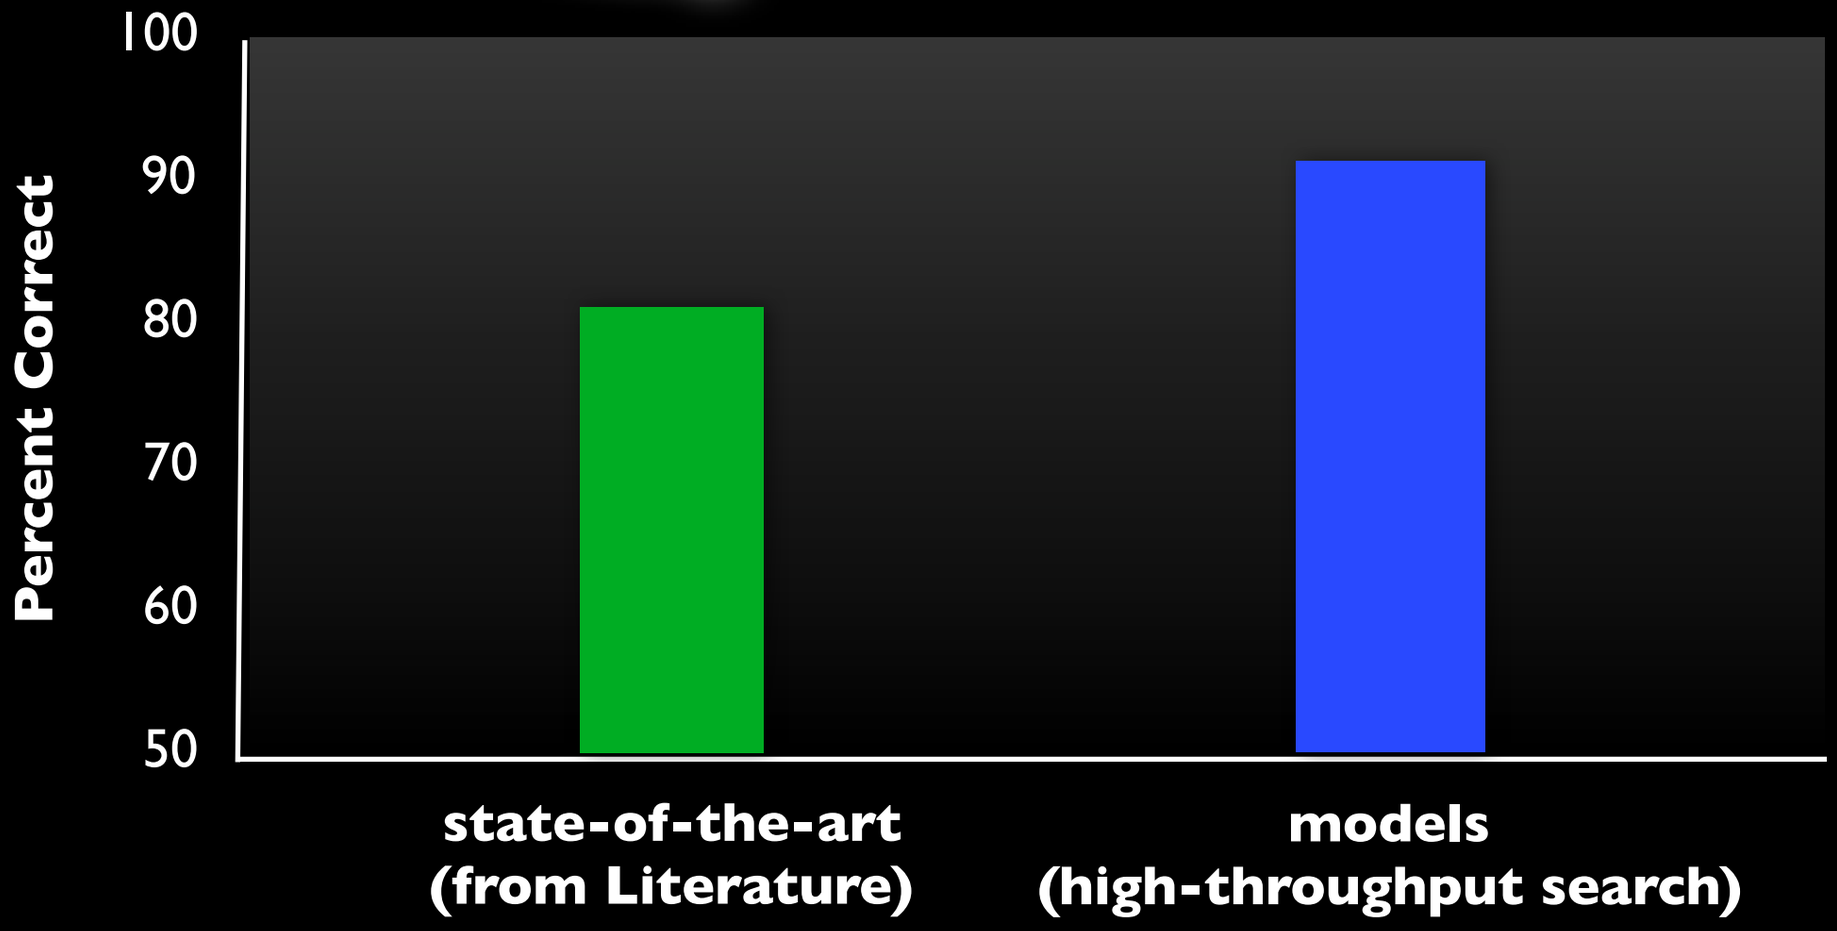

Performance

vs

vs

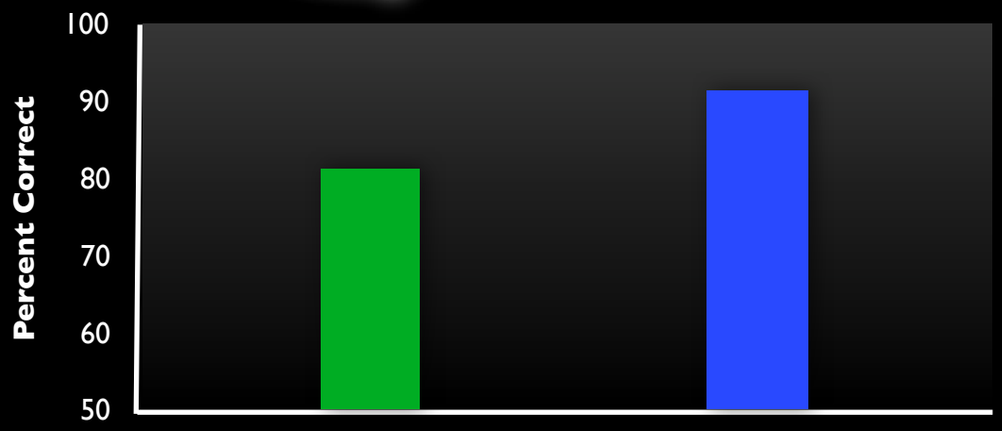

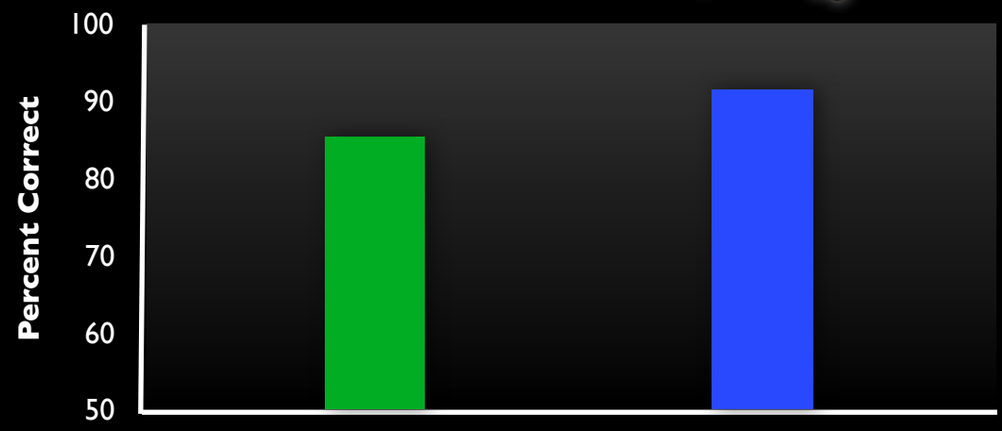

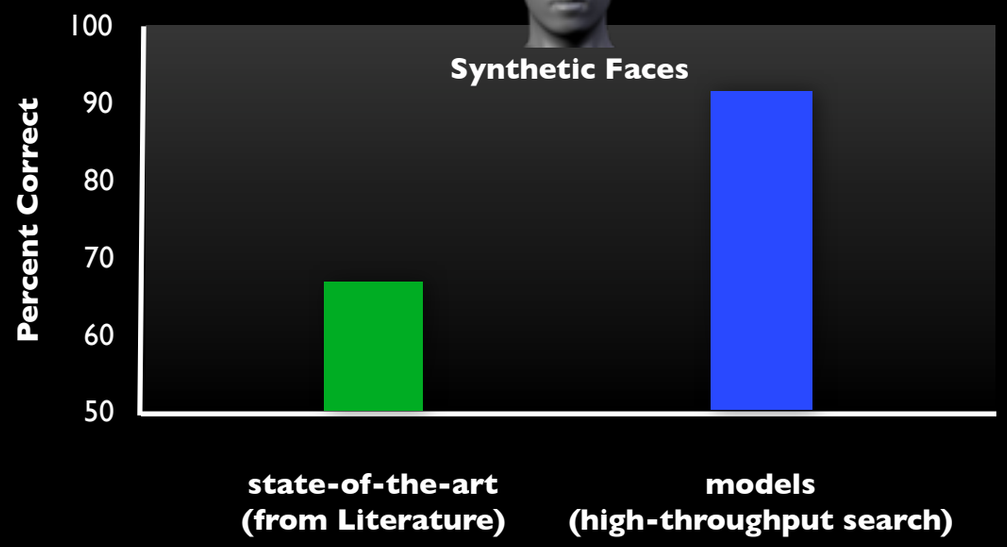

Results

3D Filterbank Convolution

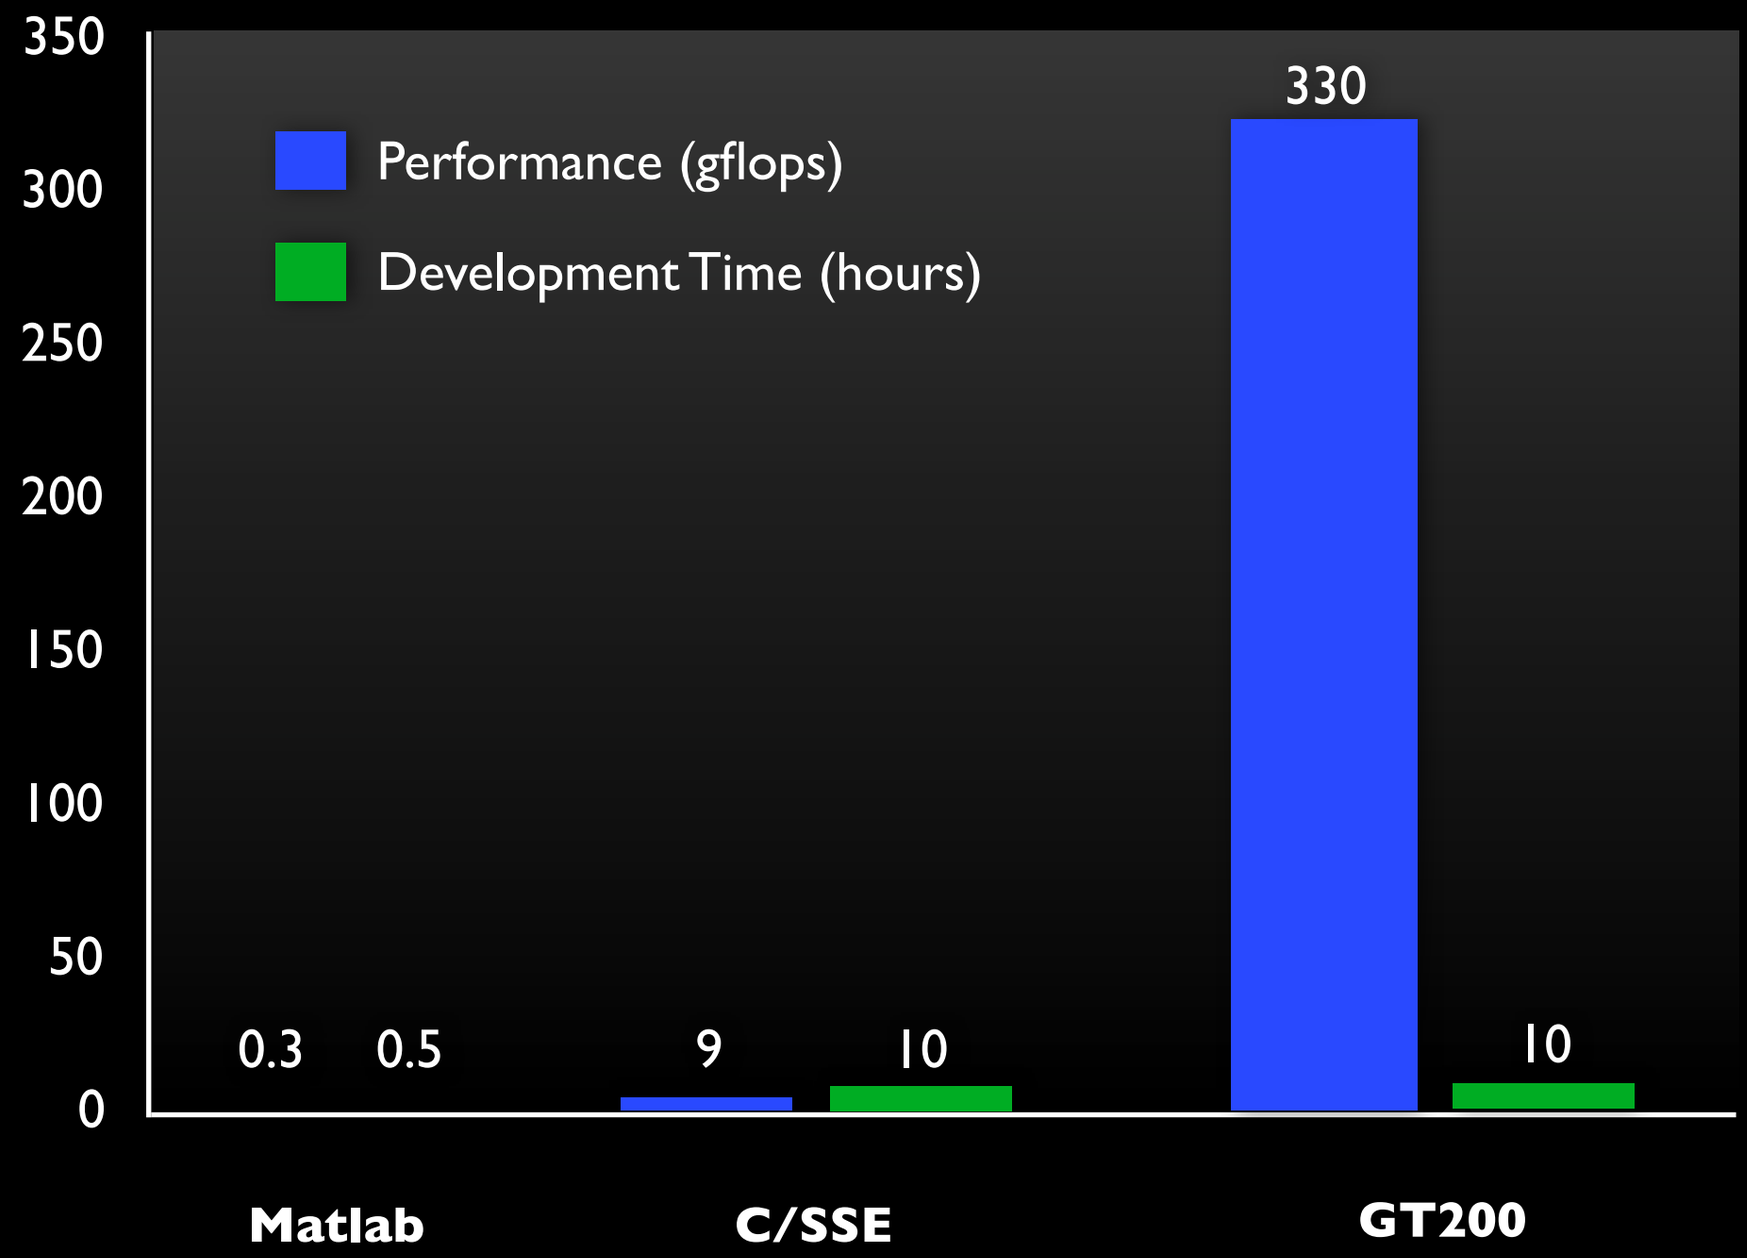

More Information

Nicolas Pinto
Unlocking Biologically-Inspired Computer Vision: a High-Throughput Approach
Session Id#1434
Thursday 5 pm Garden

Multiscale Hemodynamics

Massimo Bernaschi - Michelle Borkin - Ahmet Coskun - Charles Feldman - Efthimios Kaxiras - Simone Melchionna - Dimitris Mitsouras - Hanspeter Pfister - Frank Rybicki - Joy Sircar - Michael Steigner - Peter Stone - Sauro Succi - Frederick Welt - Amanda Whitmore

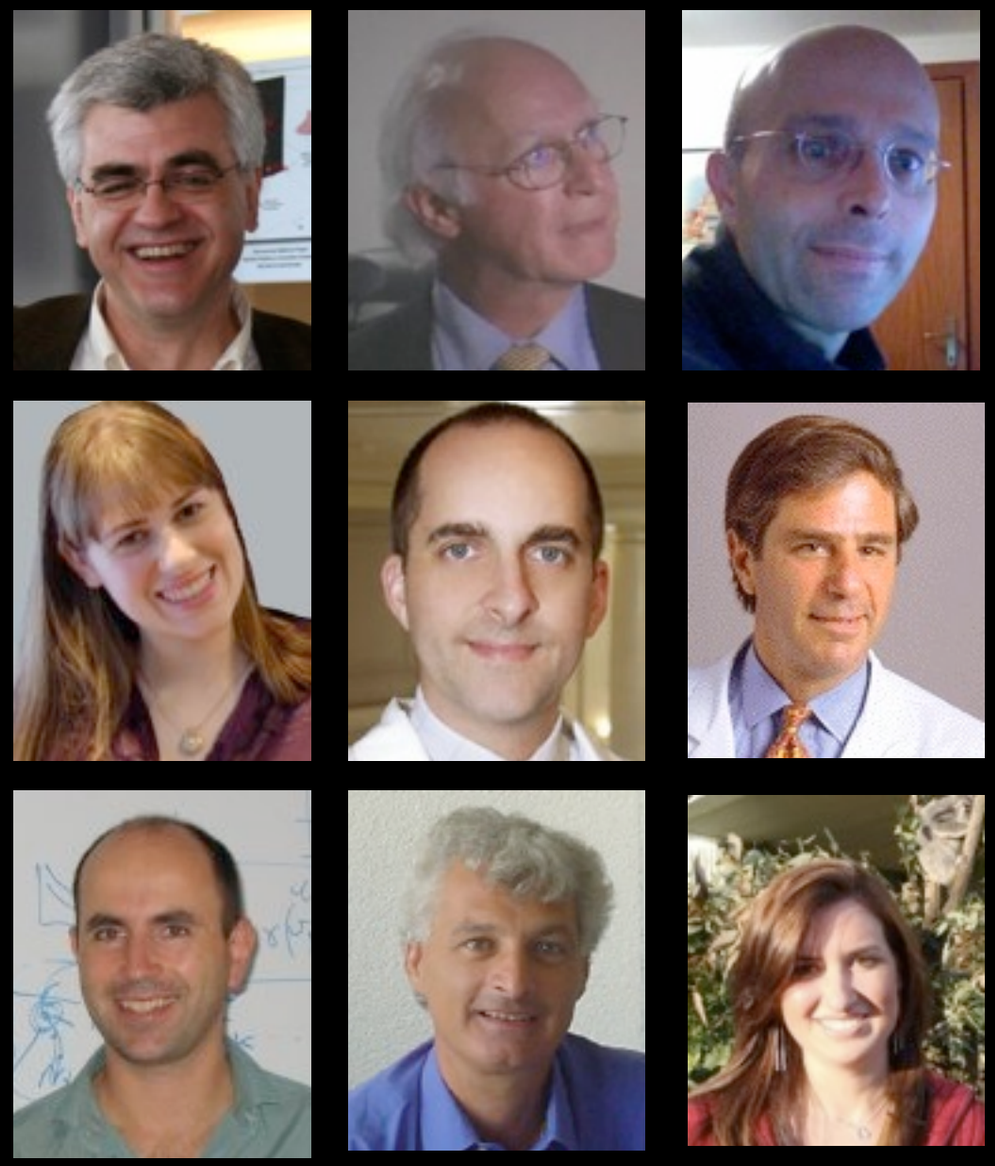

The Scientific Challenge

- Predict and prevent heart attacks

Hemodynamics Pipeline

CT Patient Data

Hemodynamics Pipeline

CT Patient Data → Reconstruction

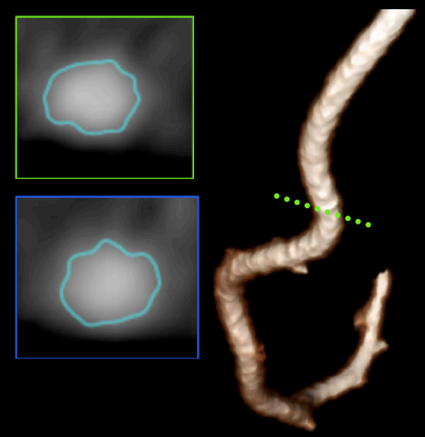

Hemodynamics Pipeline

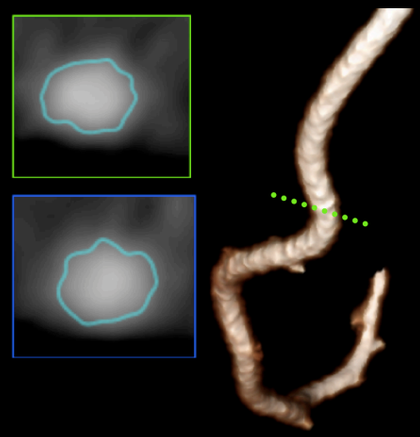

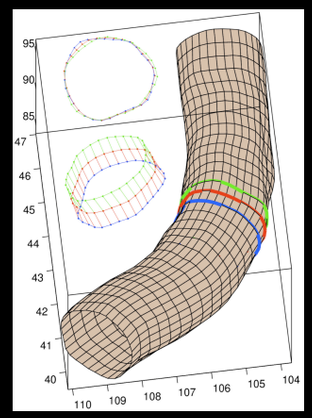

CT Patient Data → Reconstruction → Post-Processing

Hemodynamics Pipeline

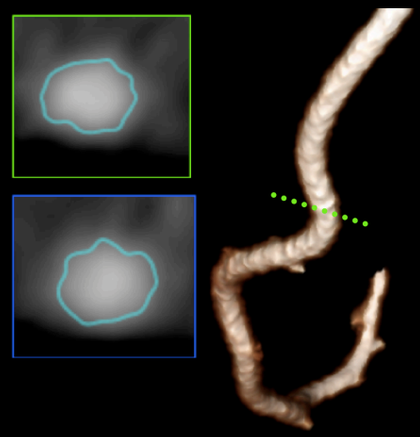

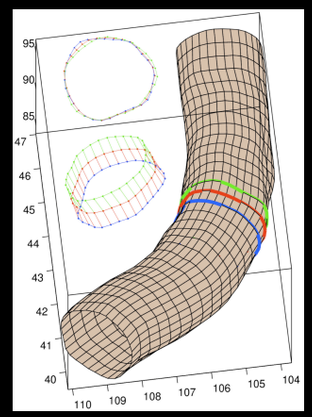

CT Patient Data → Reconstruction → Post-Processing

↓

Lattice Boltzman & Molecular Dynamics

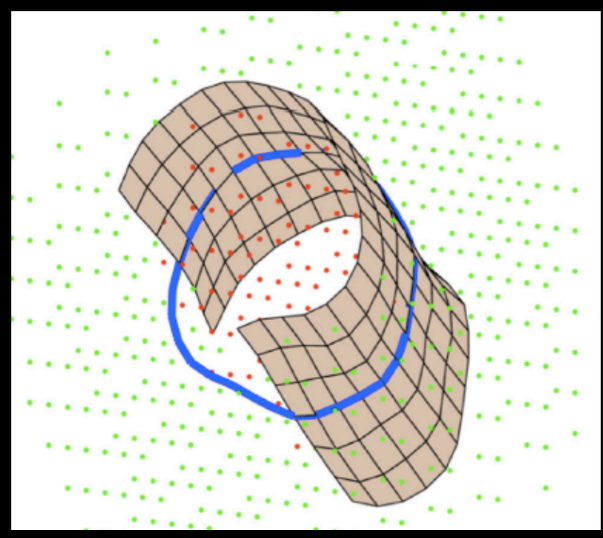

Hemodynamics Pipeline

CT Patient Data

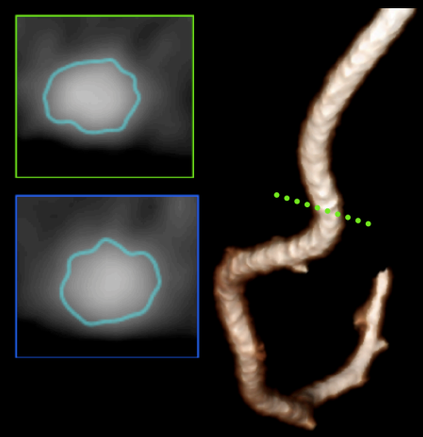

Reconstruction

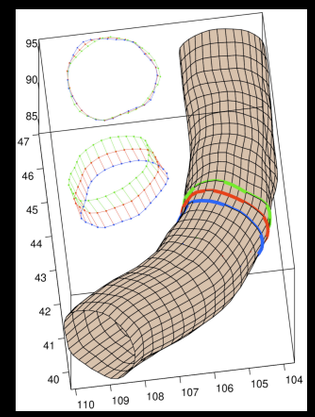

Post-Processing

Lattice Boltzman & Molecular Dynamics

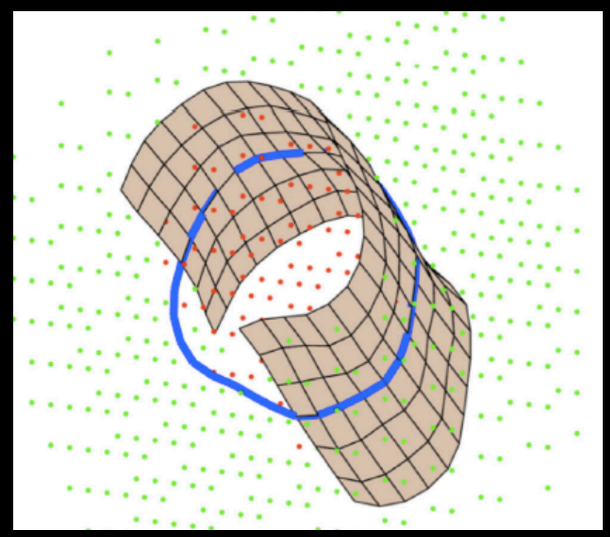

Visualization & Analysis

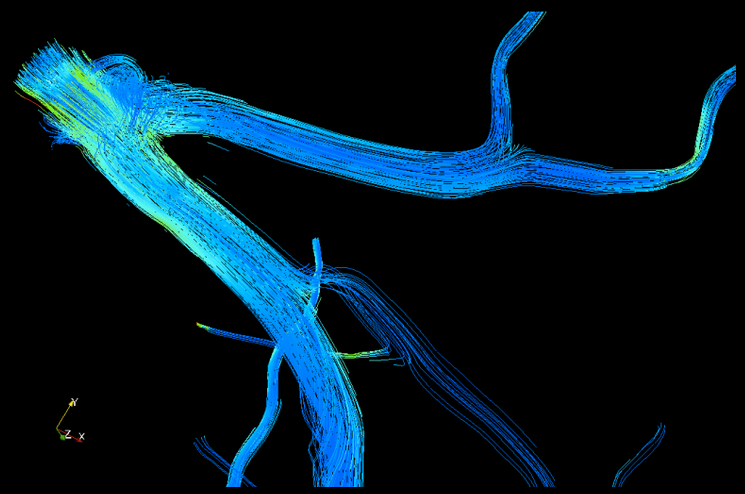

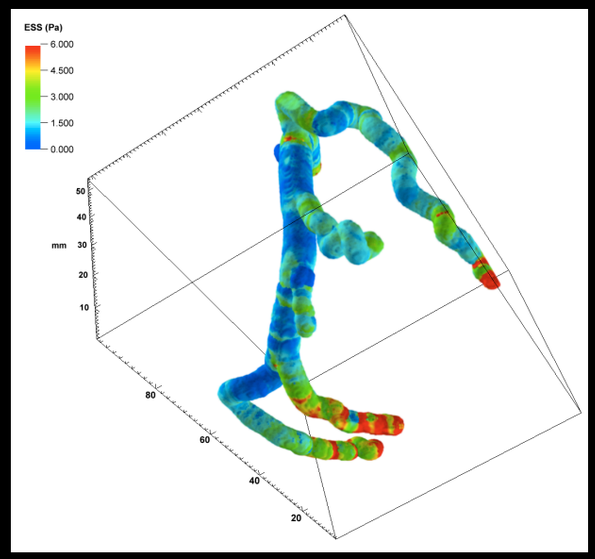

Multiscale Hemodynamics

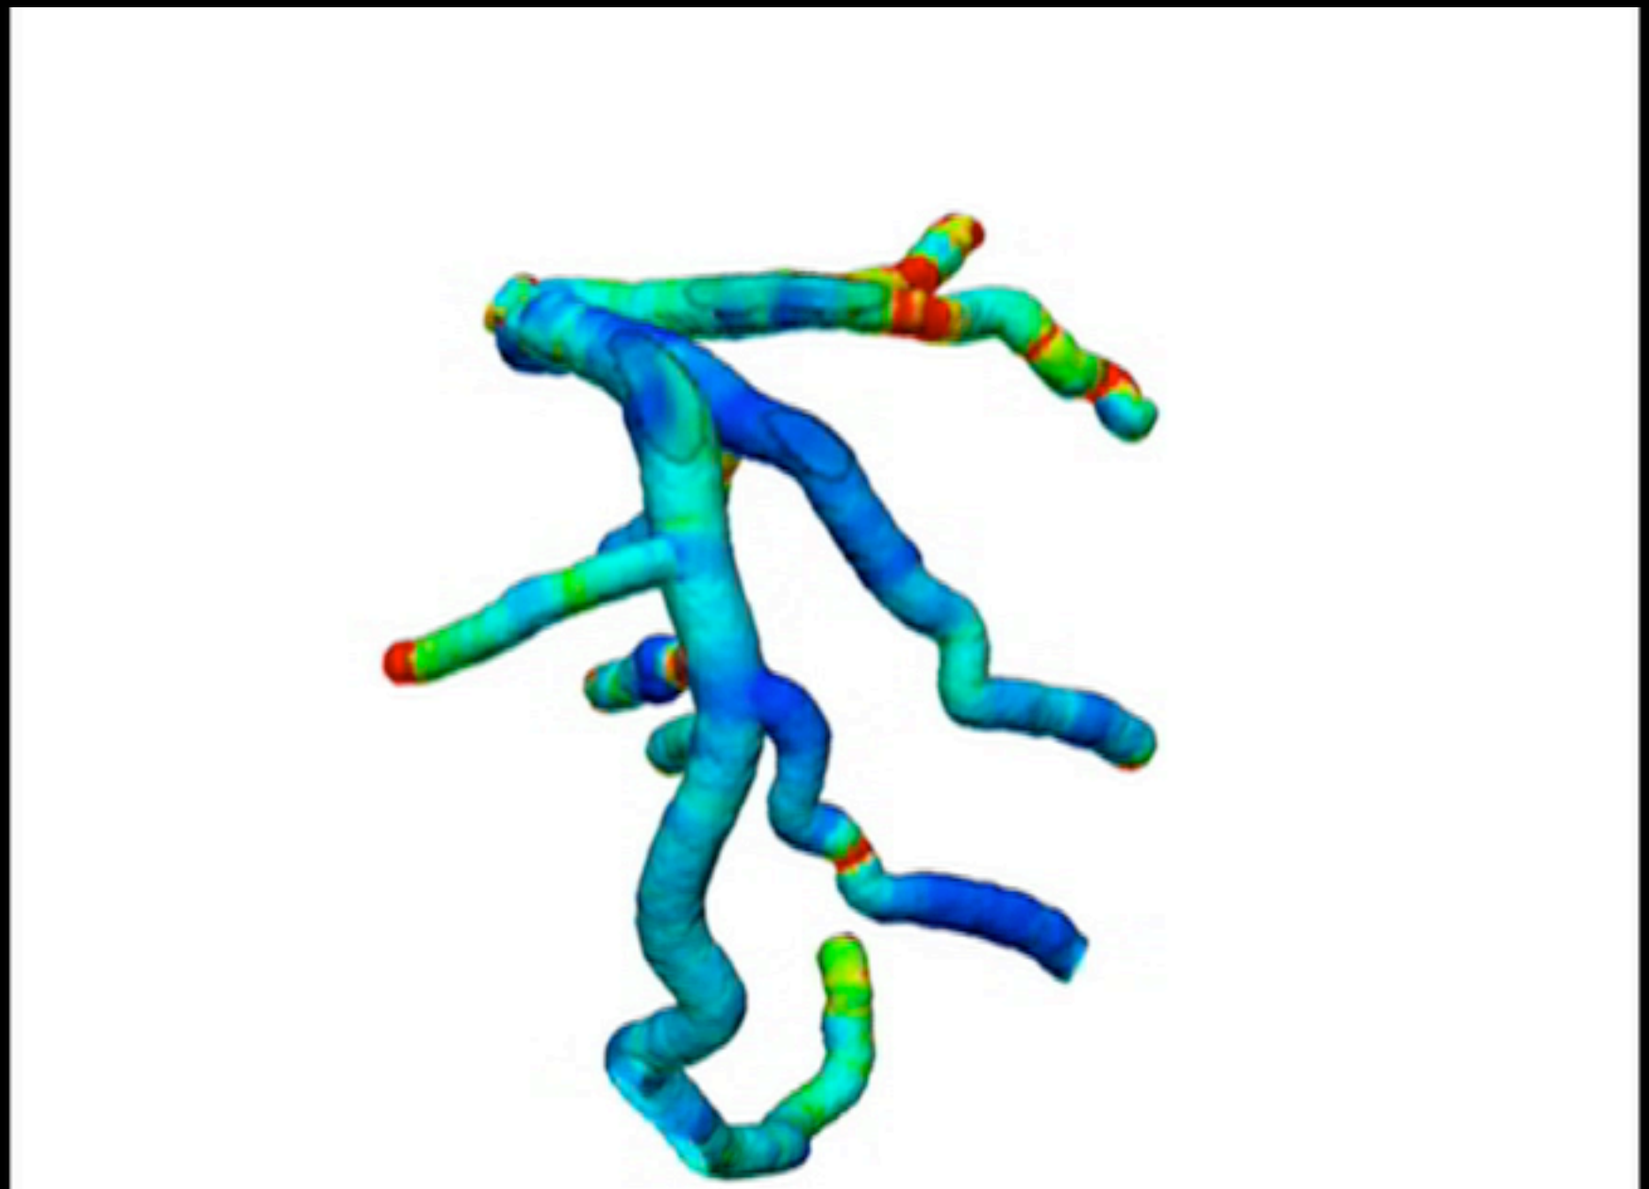

Results

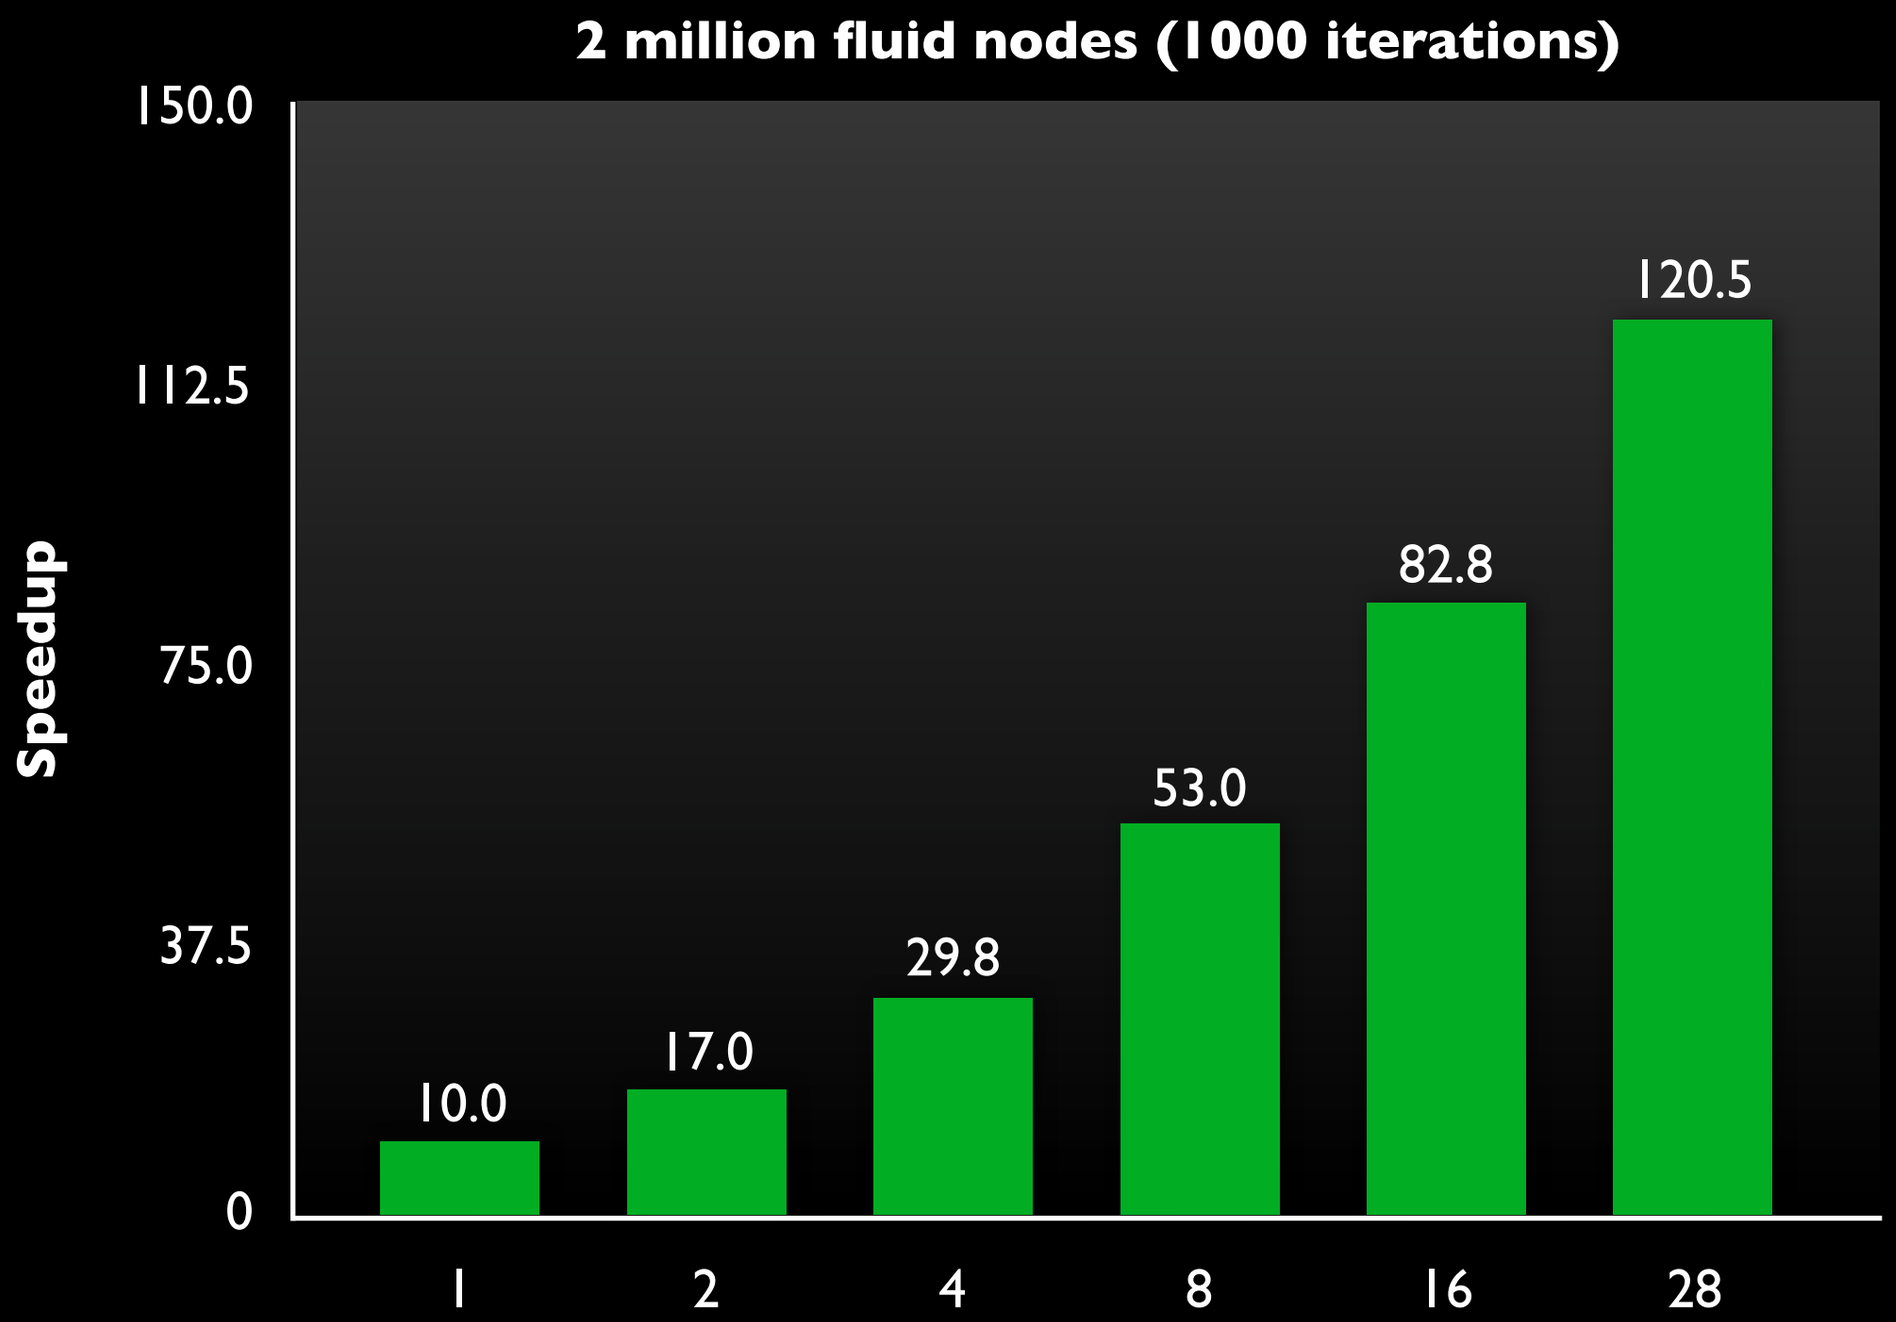
2 million fluid nodes (1000 iterations)
Speedup
150.0
112.5
75.0
37.5
0
10.0
17.0
29.8
53.0
82.8
120.5
1
2
4
8
16
28
Number of C1060 NVIDIA GPUs

The Future

- Minimally invasive early detection

UCLA Medical Center Newsletter

HTC Lessons

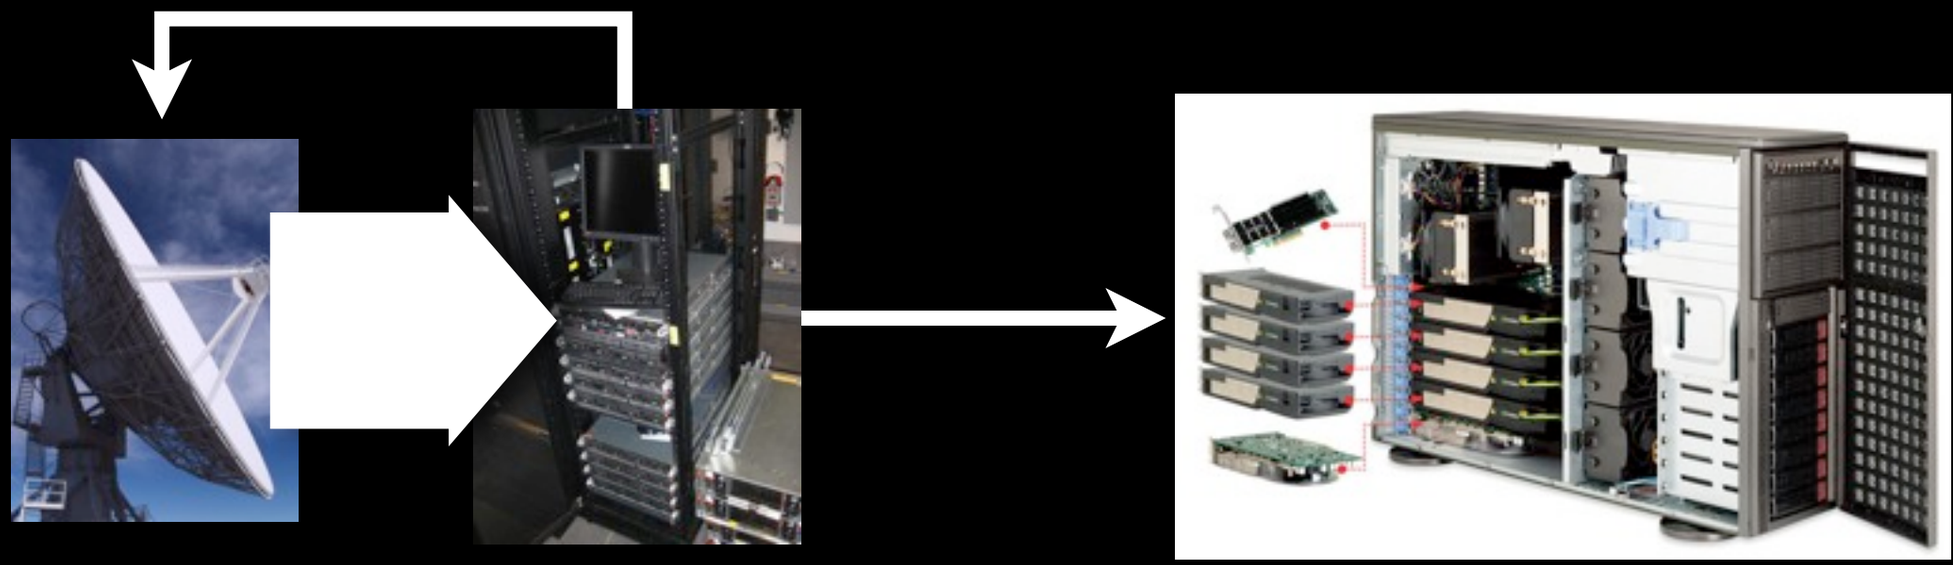

HTC Lessons

- Move the computation to the source

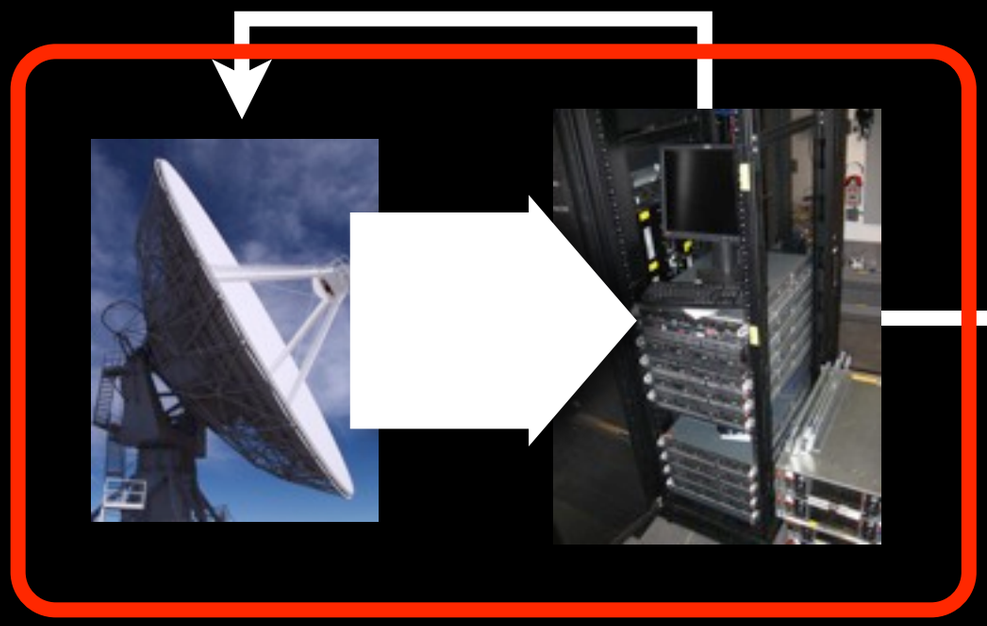

HTC Lessons

- Quickly reduce the data to its essentials

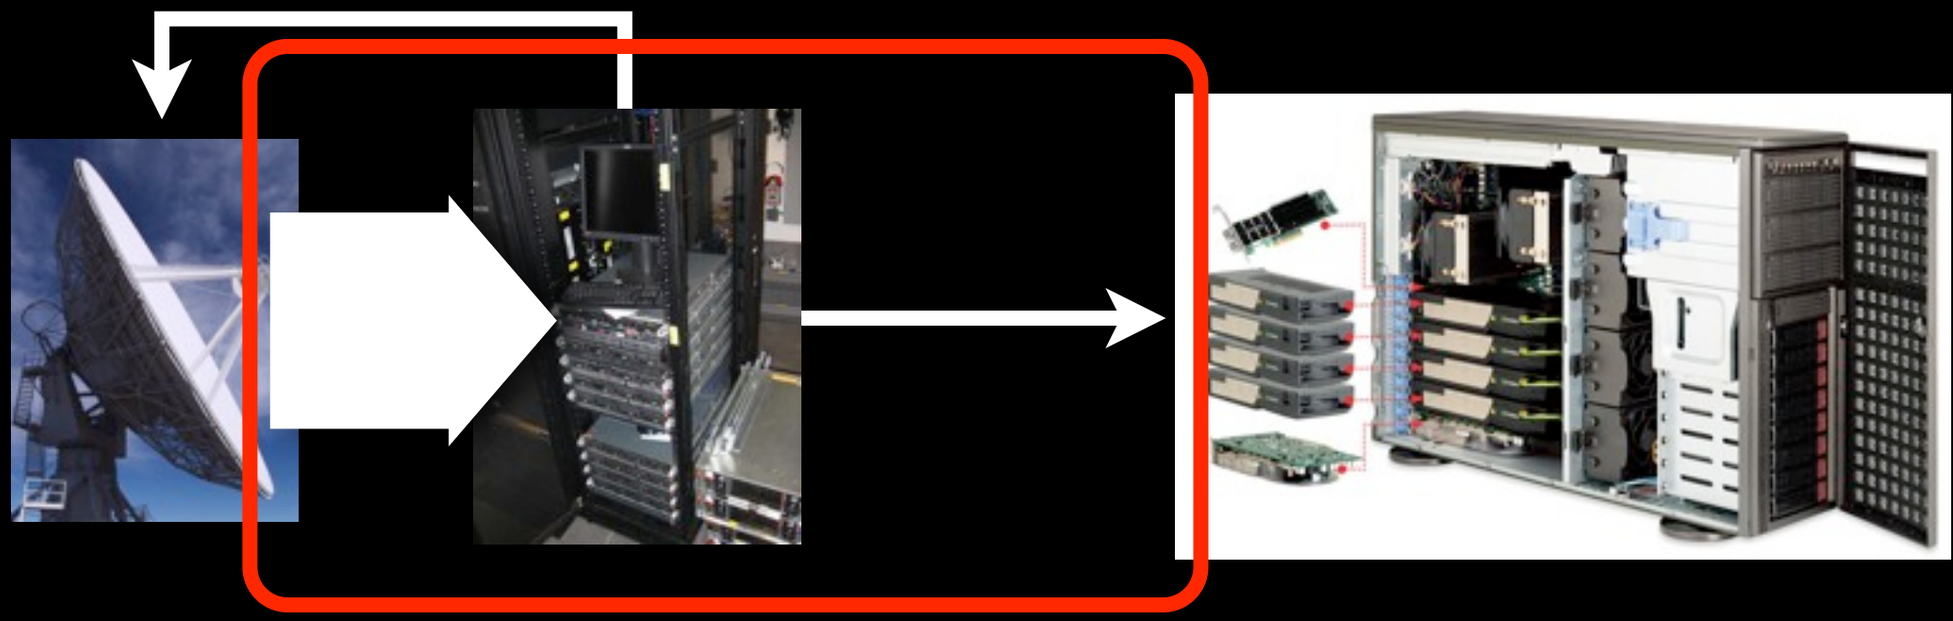

HTC Lessons

- Use feedback to optimize the acquisition

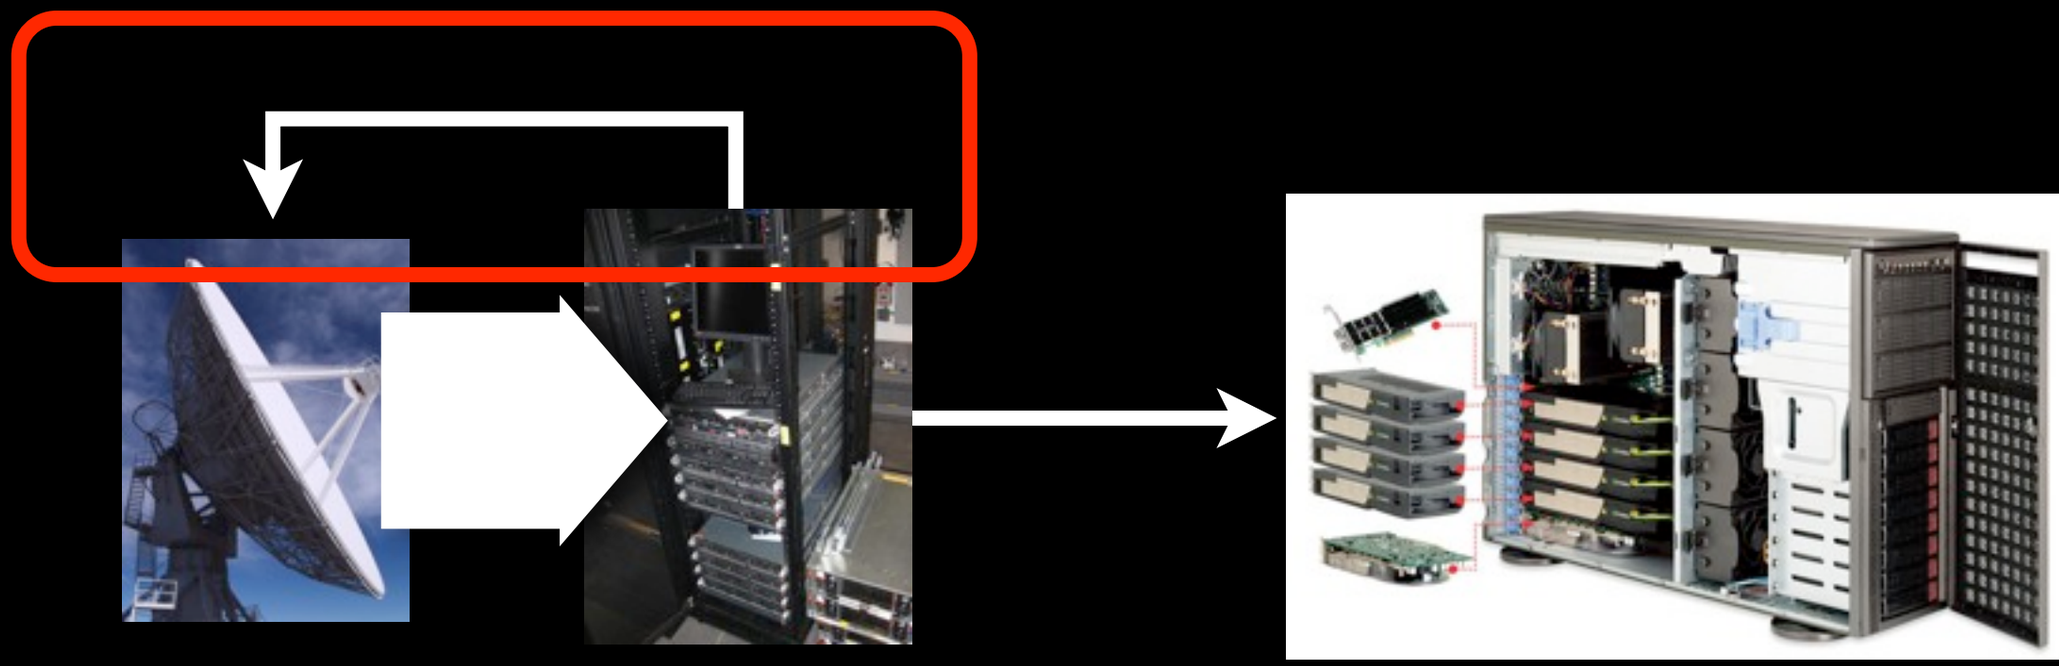

HTC Lessons

- Provides scalable and power efficient systems

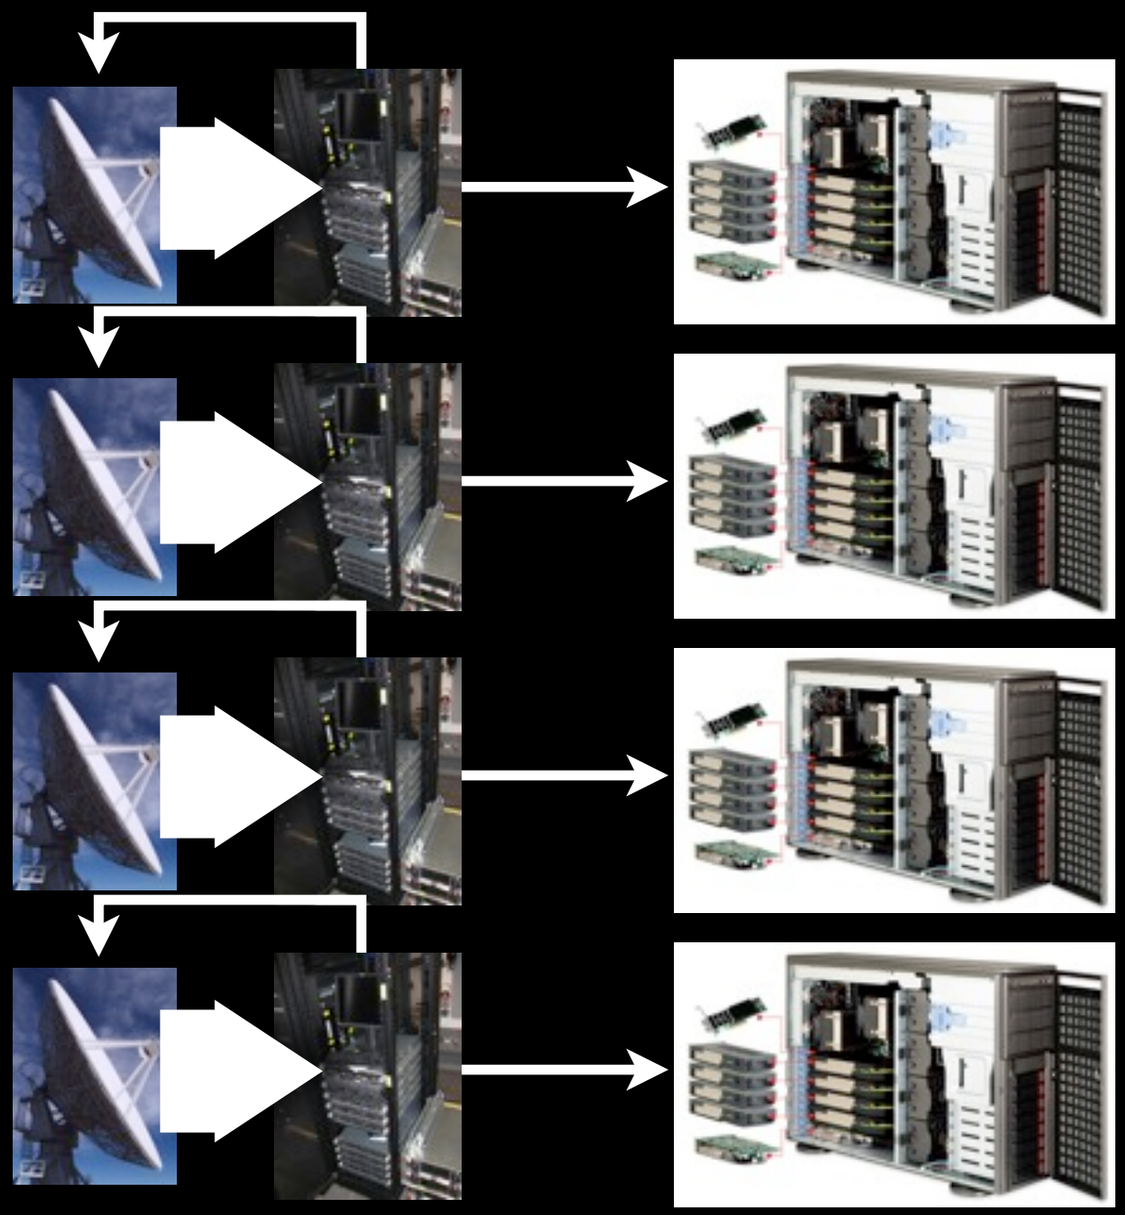

Higher-Level Programming Models

Higher-Level Programming Models

- Domain Specific Languages (DSLs)

Higher-Level Programming Models

- Domain Specific Languages (DSLs)
- E.g., Tensor Contraction Engine (TCE)

$$T1_{bcdf} = \sum_{el} B_{befl} \times D_{cdel}$$
$$T2_{bcjk} = \sum_{df} T1_{bcdf} \times C_{dfjk}$$
$$S_{abij} = \sum_{ck} T2_{bcjk} \times A_{acik}$$

(a) Formula sequence

```
T1=0; T2=0; S=0;
for b, c, d, e, f, l
  [ T1bcdf += Bbefl Dcdel
for b, c, d, f, j, k
  [ T2bcjk += T1bcdf Cdfjk
for a, b, c, i, j, k
  [ Sabij += T2bcjk Aacik
```

(b) Direct implementation (unfused code)

```
S = 0;
for b, c
  [ T1f = 0; T2f = 0;
    for d, f
      [ for e, l
          [ T1f += Bbefl Dcdel
        for j, k
          [ T2fjk += T1f Cdfjk
    for a, i, j, k
      [ Sabij += T2fjk Aacik
```

(c) Memory-reduced implementation (fused)

[Baumgartner et al. 05]

Scalable Programming Models

Plug & Play Parallel High-Throughput I/O

- Serial:
  - USB 3.0 (5 Gbp s-1)?
  - Intel Light Peak (100 Gbp s^{-1})?
- Parallel?

HTC Appliances

HTC Appliances

Surveillance

Cars

Broadcast

Video Conferencing

Chip Sensors

Gene Array Chips

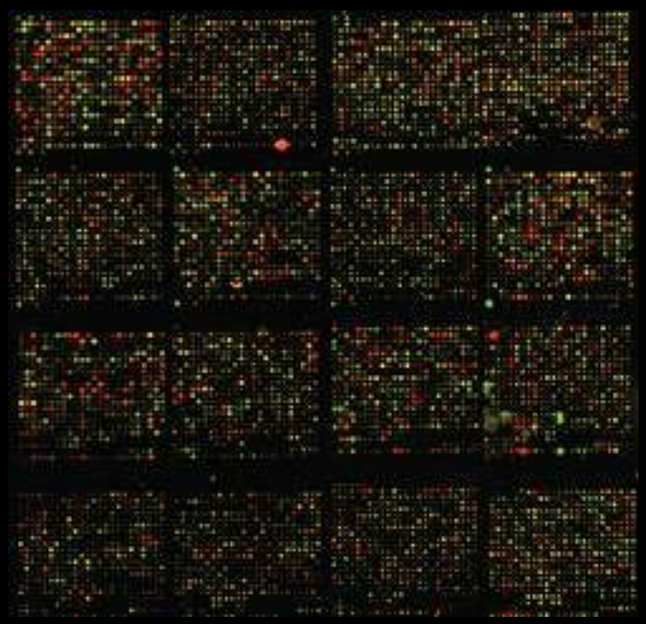

Chip-NMR Biosensor

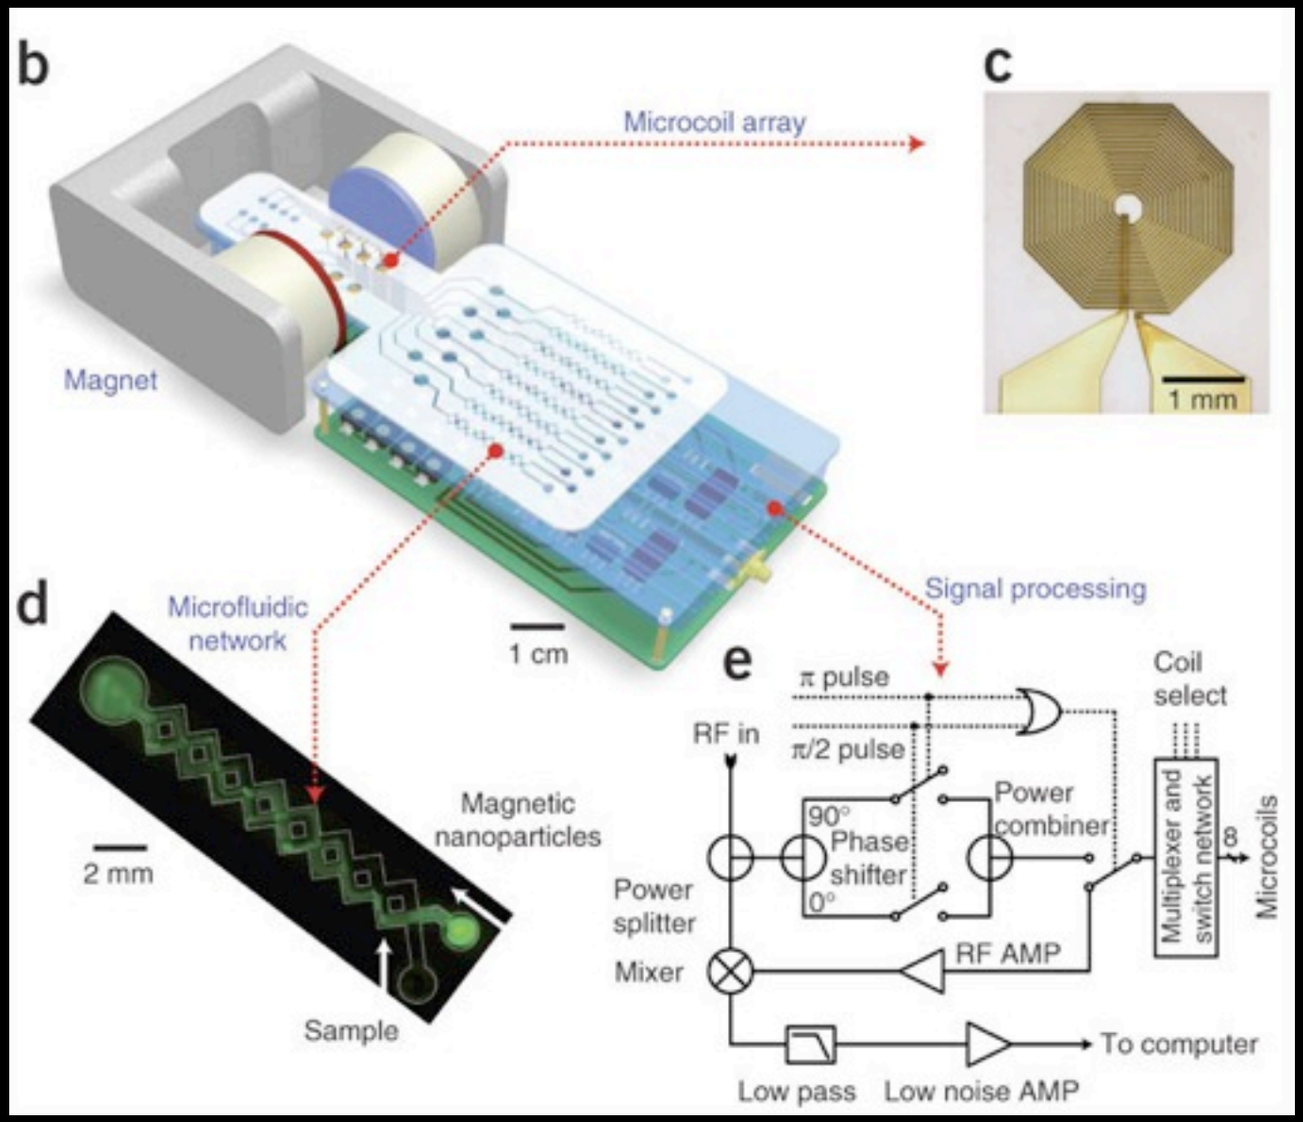

Hakho Lee, Eric Sun, Donhee Ham & Ralph Weissleder, SEAS & HMS

Tricorders

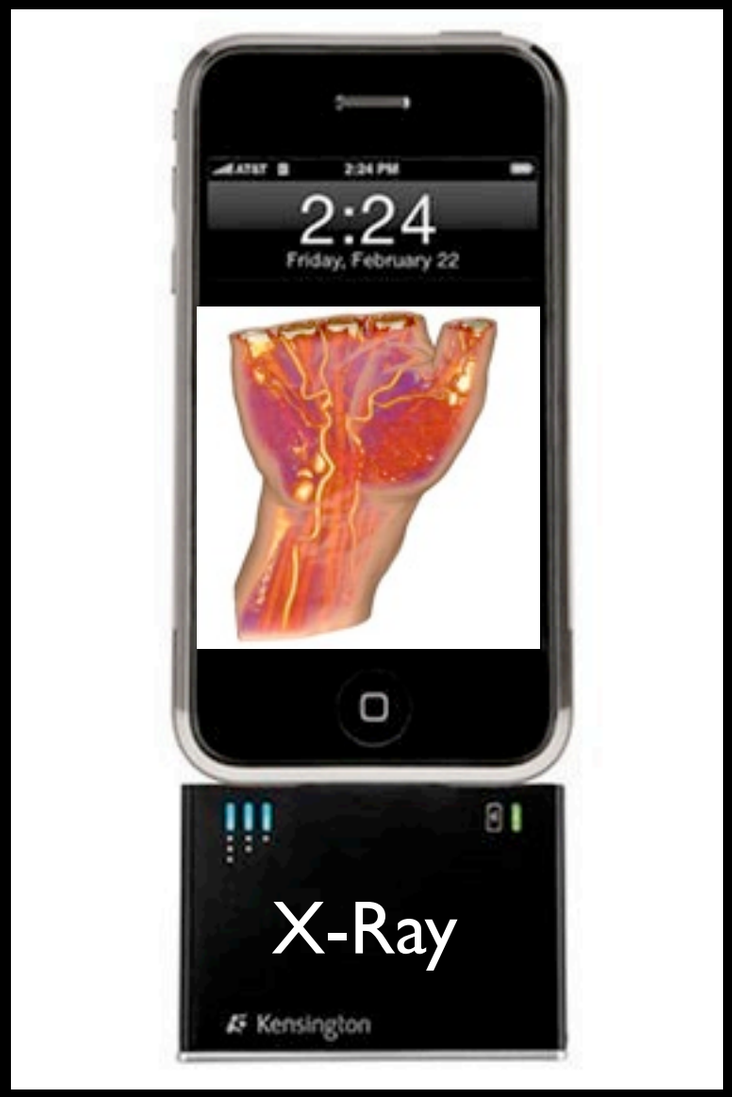

21st Century

23rd Century

Acknowledgements

- NVIDIA
- Microsoft
- NSF CDI PHY-0835713, Austrian Research Promotion Agency FFG, Vienna Science and Technology Fund WWTF